

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**

THIS PAGE BLANK (USPTO)

PCT

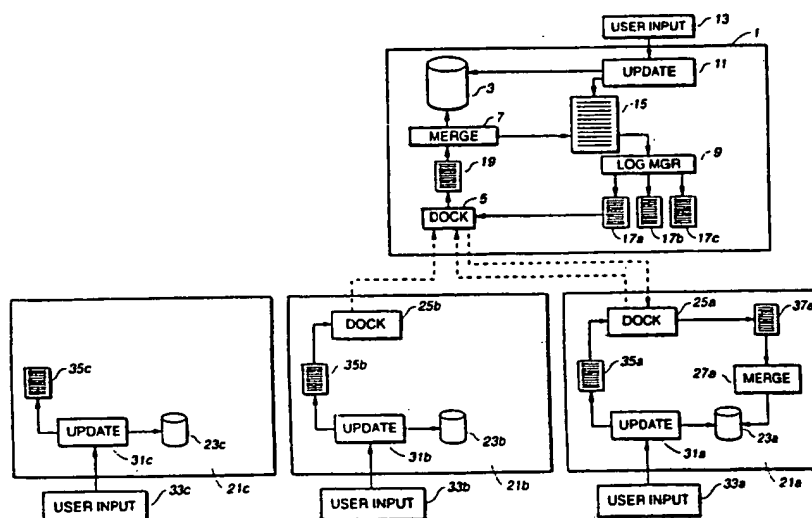
WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 15/16, 17/30, 17/40		A3	(11) International Publication Number: WO 98/38762
			(43) International Publication Date: 3 September 1998 (03.09.98)
(21) International Application Number: PCT/US98/02756 (22) International Filing Date: 24 February 1998 (24.02.98) (30) Priority Data: 60/039,167 26 February 1997 (26.02.97) US (63) Related by Continuation (CON) or Continuation-in-Part (CIP) to Earlier Application US 60/039,167 (CON) Filed on 26 February 1997 (26.02.97) (71) Applicant (for all designated States except US): SIEBEL SYSTEMS, INC. [US/US]; 1855 South Grant Street, San Mateo, CA 94402 (US). (72) Inventors; and (75) Inventors/Applicants (for US only): ROTHWEIN, Thomas, Michael [US/US]; 314 Ballymore Circle, San Jose, CA 95136 (US). COKER, John, L. [US/US]; 723 Chateau Drive, Hillsborough, CA 94010 (US). (74) Agents: GOLDMAN, Richard, M.; Cooley Godward LLP, 3000 El Camino Real, Five Palo Alto Square, Palo Alto, CA 94306-2155 (US) et al.			(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, GW, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG). Published With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments. (88) Date of publication of the international search report: 10 December 1998 (10.12.98)

(54) Title: DETERMINING VISIBILITY TO A REMOTE DATABASE CLIENT



(57) Abstract

A method and system for establishing and maintaining a secure TCP/IP session between a server (1) having a database (3) and a client (21-a, 21-b, 21-c), the method includes sending a hello message from the client to the server (1) to establish client authorization, the server (1) sends a server session ID and a session private key to the client to establish the session.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US98/02756

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : G06F 15/16, 17/30, 17/40

US CL : 395/200.47, 200.49, 200.33, 200.67

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 395/200.47, 200.49, 200.33, 200.67

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
APS, INTERNET

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5,590,199 A (KRAJEWSKI, JR. et al) 31 December 1996, col. 5 lines 14-68 and column 6 line 1 et seq.	1-16 and 18-27
X,P	US 5,657,390 A (ELGAMAL et al) 12 August 1997, col. 19 line 22 et seq col.25 and columns 29-31.	1-16 and 18-27
Y	US 5,319,710 A (ATALLA et al) 07 June 1994, col. 2 lines 10-68 and columns 3-5.	1-16 and 18-27
Y	ATKINS ET AL, PGP Message Exchange Formats RFC 1991, August 1996	1-16 and 18-27



Further documents are listed in the continuation of Box C.



See patent family annex.

•

Special categories of cited documents:

A

document defining the general state of the art which is not considered to be of particular relevance

E

earlier document published on or after the international filing date

L

document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

O

document referring to an oral disclosure, use, exhibition or other means

P

document published prior to the international filing date but later than the priority date claimed

T

later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X

document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y

document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

A

document member of the same patent family

Date of the actual completion of the international search

17 JUNE 1998

Date of mailing of the international search report

13 OCT 1998

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

MEHMET GECKIL

Telephone No. (703) 305-9676

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US98/02756

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	DCE 1.1 Online Documentation Transarc Corporation 1996 pp. 1-34 http://www.ux1.eiu.edu/~csjay/dce/intro_to_dce_4.html	1-16 and 18-27

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US98/02756

Box I Observations where certain claims were found unsearchable (Continuation of Item 1 of first sheet)

This international report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☒ Claims Nos.: 17
because they relate to subject matter not required to be searched by this Authority, namely:

Claim recite non functional material. Also the numbers recited in the claims e.g., 301 and 311 does not exist in the figures or in the specification.
2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box II Observations where unity of invention is lacking (Continuation of item 2 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest.
☐ No protest accompanied the payment of additional search fees.

BEST AVAILABLE COPY

BEST AVAILABLE COPY

Continuation

Lead Article

led by this An

claims eq

at comply
not, specific

with the

of items

information

icant, this

ing an addi

ed by the

and

icant.

acted by

accompanied by

ment of addition

PCTWORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04L	A2	(11) International Publication Number: WO 98/38762 (43) International Publication Date: 3 September 1998 (03.09.98)
<p>(21) International Application Number: PCT/US98/02756</p> <p>(22) International Filing Date: 24 February 1998 (24.02.98)</p> <p>(30) Priority Data: 60/039,167 26 February 1997 (26.02.97) US</p> <p>(63) Related by Continuation (CON) or Continuation-in-Part (CIP) to Earlier Application US 60/039,167 (CON) Filed on 26 February 1997 (26.02.97)</p> <p>(71) Applicant (for all designated States except US): SIEBEL SYSTEMS, INC. [US/US]; 1855 South Grant Street, San Mateo, CA 94402 (US).</p> <p>(72) Inventors; and (75) Inventors/Applicants (for US only): ROTHWEIN, Thomas, Michael [US/US]; 314 Ballymore Circle, San Jose, CA 95136 (US). COKER, John, L. [US/US]; 723 Chateau Drive, Hillsborough, CA 94010 (US).</p> <p>(74) Agents: GOLDMAN, Richard, M.; Cooley Godward LLP, 3000 El Camino Real, Five Palo Alto Square, Palo Alto, CA 94306-2155 (US) et al.</p>		<p>(81) Designated States: -AL- -AM- -AT, AU, -AZ- -BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, GW, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).</p> <p>Published <i>Without international search report and to be republished upon receipt of that report.</i></p>
<p>(54) Title: METHOD OF DETERMINING VISIBILITY TO A REMOTE DATABASE CLIENT OF A PLURALITY OF DATABASE TRANSACTIONS USING A NETWORKED PROXY SERVER</p>		
<p>(57) Abstract</p> <p>A method of establishing and maintaining a secure TCP/IP session between a server (301) having a database and a client (311). The method includes sending a hello message from the client (311) to the server (301) to establish client authorization, the server (301) sends a server session ID and a session private key to the client (311) to establish the session. Also disclosed is an article of manufacture that includes a computer usable medium having computer readable program code for causing the above method, and a program storage device tangibly embodying program code to effect the above method.</p>		

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon	KR	Republic of Korea	PL	Poland		
CN	China	KZ	Kazakhstan	PT	Portugal		
CU	Cuba	LC	Saint Lucia	RO	Romania		
CZ	Czech Republic	LI	Liechtenstein	RU	Russian Federation		
DE	Germany	LK	Sri Lanka	SD	Sudan		
DK	Denmark	LR	Liberia	SE	Sweden		
EE	Estonia			SG	Singapore		

**METHOD OF DETERMINING THE VISIBILITY TO A REMOTE
DATABASE CLIENT OF A PLURALITY OF DATABASE
5 TRANSACTIONS USING A NETWORKED PROXY SERVER**

INTRODUCTION

Technical Field

10 This invention relates to a system and method for providing updates to a network of partially replicated relational database systems, and, more particularly, for providing efficient access to a database by a remote client using a networked proxy server.

15 **Background**

Relational databases are a commonly-employed data structure for representing data in a business or other environment. A relational database represents data in the form of a collection of two-dimensional tables. Each table comprises a series of cells arranged in rows and columns. Typically, a
20 row in a table represents a particular observation. A column represents either a data field or a pointer to a row in another table.

For example, a database describing an organizational structure may have one table to describe each position in the organization, and another table to
25 describe each employee in the organization. The employee table may include information specific to the employee, such as name, employee number, age, salary, etc. The position table may include information specific to the position, such as the position title ("salesman", "vice president", etc.), a salary range,

and the like. The tables may be related by, for example, providing in each row of the employee table a pointer to a particular row in the position table, coordinated so that, for each row in the employee table, there is a pointer to the particular row in the position table that describes that employee's position.

- 5 A relational database management system (RDBMS) supports "joining" these tables in response to a query from a user, so that the user making a query about, for example, a particular employee, may be provided with a report of the selected employee, including not only the information in the employee table, but also the information in the related position table.

10

Relational databases may be much more complex than this example, with several tables and a multiplicity of relations among them.

- 15 With the widespread use of inexpensive portable computers, it is advantageous to replicate a database onto a portable computer for reference at locations remote from the central computer. The replicated database may then be referenced by the user of the portable computer, without requiring reference to the main database, which may be maintained at a central location inconvenient to the user of the portable computer. However, there are a number of difficulties with the use of a replicated database.
- 20

- One disadvantage is that a full copy of the central database may require more data storage than is desired or economical. For example, a salesman working in the field may need to refer to the database for information regarding sales opportunities in his sales area, but have no need to refer to any information regarding sales opportunities outside of his area. One possible approach to reduce the amount of required data storage is to simply replicate only that portion of the database that is needed by the user. However, this approach does not recognize that the criteria to determine which portions of the data are required is likely to vary over time. For example, the salesman may
- 25
- 30

have a new city added to his territory. Under conventional approaches, the salesman would need to re-replicate his local copy of the database, this time selecting data including the added city. Such a practice is inconvenient, subject to error, and time-consuming.

5

A further disadvantage to a replicated database is the difficulties encountered in attempting to update data using the replicated copy. A change made to the replicated database is not made to the central database, leading to a discrepancy between the information that is stored in the replicated copy of the database and the information that is stored in the central database. Although it is possible to journal modifications made to the replicated copy and apply an identical modification to the central database, one problem that this approach faces is the possibility of colliding updates; that is, where a user of a replicated copy makes a change to data that is also changed by a user of the central copy of by the user of another replicated copy.

It is therefore desirable to provide a capability to maintain one or more partially-replicated copies of a central database, in such a way that the degree of replication may be easily changed without requiring a refresh of the entire replicated database, and that permits updates to be coordinated among users of the central database and users of the partially replicated databases. In addition, it is also desirable to provide access to the central database by the users of the partially replicated databases over a wide area network, so that it is not necessary for each user to call the central database directly. The ability to use the Internet, a commercial online provider's network, or even a corporate intranet, significantly increases the utility of the partially replicated databases by providing excellent connectivity at relatively minimal cost. Further, the software and the replicated databases may be integrated with other services which use the Internet.

30

SUMMARY OF THE INVENTION

The present invention is directed to a method of maintaining a partially replicated database in such a way that updates made to a central database, or to another partially replicated database, are selectively propagated to the partially replicated database. Updates are propagated to a partially replicated database if the owner of the partially replicated database is deemed to have visibility to the data being updated. Visibility is determined by use of predetermined rules stored in a rules database. In one aspect of the invention, the stored rules are assessed against data content of various tables that make up a logical entity, known as a docking object, that is being updated.

In another aspect of the invention, the stored rules are assessed against data content of one or more docking objects that are not necessarily updated, but that are related to a docking object being updated. In one embodiment, the visibility attributes of the related docking objects are recursively determined.

In yet another aspect of the invention, changes in visibility are determined to enable the central computer to direct the nodes to insert the docking object into its partially replicated database. Such changes in visibility are determined so as to enable the central computer to direct a node to remove a docking object from its partially replicated database.

In a further aspect of the invention, the predetermined rules are in declarative form and specify visibility of data based upon structure of the data without reference to data content.

In still another aspect of the invention, the transactions made to the database are ordered and processed in such a way as to reduce the computational resources required to calculate the visibility of the transactions.

In another aspect of the invention, the transactions are transmitted over the Internet and provided to a networked proxy server which transmits the requests to the central computer for database access.

5 In yet another aspect of the invention, security and access control are provided for secure transmissions over an insecure network.

These and other aspects of the inventions will become apparent to one skilled in the art by reference to the following drawings and detailed
10 description.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 depicts an overview of the operation of one embodiment of the present invention.

15

Figure 2 depicts a database schema that shows the relationship of the various components that make up a Docking Object.

Figure 3 depicts steps performed by an update manager to update a
20 database.

Figure 4 depicts steps performed by a Docking Manager to transmit and/or receive one or more transaction logs.

25 Figure 5 depicts the steps performed by a merge processor to merge transaction log records into an existing database.

Figure 6 depicts the steps performed by a log manager to prepare a partial transaction log.

30

Figure 7 depicts the steps performed by a visibility calculator for calculating visibility for a docking object as invoked by a log manager.

Figure 8 depicts the steps performed to synchronize a partially replicated
5 database in response to a change in data visibility.

Figure 9 depicts the overall structure of a system embodying the present invention using a networked proxy server.

10 DESCRIPTION OF SPECIFIC EMBODIMENTS

Overview

Figure 1 depicts an overview of the operation of one embodiment of the present invention. Figure 1 depicts a central computer system 1 and three remote computer systems (or "nodes") 21-a, 21-b, and 21-c. Each of nodes 21-
15 a, 21-b and 21-c are depicted in various states of communication with central computer system 1, as will be more fully explained. Central computer system 1 includes a central database 3, a docking manager 5, a merge processor 7 and a log manager 9. Central computer system 1 additionally optionally includes update manager 11 responsive to user input 13.

20 Node 21-a is a remote computer system, such as a mobile client such as a laptop computer. Node 21-a includes a partially replicated remote database 23-a, update manager 31-a responsive to user input 33-a, docking manager 25-a and merge manager 27-a. In operation, update manager is responsive to user
25 input 33-a to make changes to remote database 23-a as directed by the operator of node 21-a. Updates made are recorded, or journaled, in node update log 35-a.

At some point at the convenience of the operator of node 21-a, node
30 docking manager 35-a is activated, and enters into communication with central

docking manager 5. Update log 35-a is taken as input by node docking manager 25-a, and provided to central docking manager 5. Central docking manager 5 creates a received node update log 19, which contains all the information that had been recorded in update log 35-a. Optionally, partial log 17-a is taken as input by central docking manager 5 and provided to node docking manager 25-a, as more fully described herein.

At some point in time, at the convenience of the operator of central computer system 1, merge processor 7 is activated. Merge processor 7 takes as input received node update log 19, and applies the updates described therein to central database 3. In the process of applying the updates from received node update log 19, merge processor journals the updates applied to central update log 15. Optionally, update manager 11, responsive to user input 12 makes additional changes to central database 3 as directed by the operator of central computer system 1. The updates made by update manager 11 are additionally journaled in central update log 15.

At some point in time, at the convenience of the operator of central computer system 1, log manager 9 is activated. Log manager 9 takes as input central update log 15 and produces as output a set of partial logs 17-a, 17-b and 17-c according to visibility rules as will be further described herein. Each of partial logs 17-a, 17-b and 17-c corresponds to one of nodes 21-a, 21-b and 21-c. When a node docking manager such as node docking manager 25-a enters into communication with central docking manager 5 and optionally requests transmission of its corresponding partial log, central docking manager 5 takes as input the appropriate partial log, such as partial log 17-a, and presents it to node docking manager 25-a. Node docking manager 25-a then replicates partial log 17-a as merge log 37-a.

At some point in the future, at the convenience of the operator of node

21-a,

merge processor 27-a is activated. Merge processor 27-a takes as input merge log 37-a, and applies the updates described therein to partially replicated database 23-a.

5

In addition to node 21-a, Figure 1 also depicts two additional nodes 21-b and 21-c. Node 21-b is depicted in communication with central computer 1. However, unlike node 21-a, the operator of node 21-b has requested only to send his updates to central computer system 1, and has not requested to be
10 presented with changes made elsewhere to be made to his partially replicated database 23-b. This may be, for example, if the operator has an urgent update that must be made as soon as possible, but does not have the time to receive updates from other nodes. Accordingly, Figure 1 shows only transmission of node update log 35-a from node docking manager 25-b to central docking
15 manager 5, and no transmission from central docking manager 5 to node docking manager 25-b. Accordingly, the merge manager for node 21-b is not activated and is not shown.

Likewise, node 21-c is depicted as not in communication with central
20 computer system 1. Accordingly, the docking manager for node 21-c is not activated and is not shown.

By the cycle described above, updates made by each of nodes 21-a, 21-b and 21-c are presented to central computer system 1, permitting central
25 database 3 to be updated accordingly. In addition, each of the updates made by each of the nodes 21-a, 21-b and 21-c, as well as updates made on central computer system 1, are routed back to each of nodes 21-a, 21-b, and 21-c, thereby keeping each of partial databases 23-a, 23-b and 23-c in
synchronization with each other and with central database 3.

30

Database Structure

The synchronization of central database 3 with node databases 23-a, 23-b and 23-c is performed using a construct called a Docking Object. A Docking Object consists of Member Tables (including one Primary Table), Visibility Rules, Visibility Events, and related Docking Objects.

A Member Table is a table of the relational database that makes up a docking object. When a docking object is propagated from central database 3 to one of node databases 23-a, 23-b or 23-c, the propagation takes the form of an insertion into each of the Member Tables associated with the particular docking object. Similarly, when a docking object is scheduled to be removed from a database, that removal consists of deleting records from the member tables associated with the docking object. For example, a docking object that represents a sales opportunity may include tables that represent the opportunity itself (e.g., named "S_OPTY"), the product whose sale is represented by the opportunity (e.g., named "S_OPTY_PROD"), the contact for the opportunity (e.g., named "S_OPTY_CONTACT"), etc. Each of these tables is said to be a member table of the "Opportunity Docking Object."

A Primary Table is a Member Table that controls whether a particular instance of a Docking Object is visible to a particular node. The Primary Table has a Primary Row-ID value that is used to identify a row of the Primary Table being updated, deleted or inserted. For example, the "Opportunity Docking Object" may have as a primary table the table S_OPTY. The row-id of that table, i.e., S_OPTY.row_id, is the Primary Row-ID for the Opportunity Docking Object.

A Visibility Rule is a criterion that determines whether a particular instance of a Docking Object is "visible" to a particular node. If a Docking Object is visible to a particular node, that node will receive updates for data in

the Docking Object. Visibility Rules are of two types, depending on the field RULE_TYPE. A Visibility Rule with a RULE_TYPE of "R" is referred to as an SQL Rule. An SQL Rule includes a set of Structured Query Language (SQL) statements that is evaluated to determine if any data meeting the criteria specified in the SQL statements exists in the Docking Object. If so, the Docking Object is visible to the node. A Visibility Rule with a RULE_TYPE of "O" is referred to as a Docking Object Rule. A Docking Object Rule specifies another Docking Object to be queried for visibility. If the specified Docking Object is visible, then the Docking Object pointing to it is also visible.

10

A Related Docking Object is a Docking Object that is propagated or deleted when the Docking Object under consideration is propagated or deleted. For example, an Opportunity Docking Object may have related Docking Objects representing the sales contacts, the organizations, the products to be sold, and the activities needed to pursue the opportunity. When an Opportunity Docking Object is propagated from Central Database 3 to one of node databases 23, the related docking objects are also propagated.

Figure 2 depicts a database schema that shows the relationship of the various components that make up a Docking Object. The schema is a meta-database, in that it does not describe the data being accessed in the database. Rather, the schema is a separate database that defines the structure of the database being accessed. That is, it is a database comprising tables that describe the relationships and data contexts of another database.

Each of the tables shown in Figure 2 is a table in a relational database, and as such is in row-column form. Many columns represent fields that are common to all the illustrated tables. Such fields include for example, a ROW_ID to identify a particular row in the table, as well as fields to tack the date and time that a row was created and last modified, and the identity of the user who created or modified the row. In addition, each table contains fields

specific to that table, and which are described in detail below.

Table S_DOBJ 61 describes the Docking Objects in an application.

Table S_DOBJ 61 includes the fields OBJ_NAME and PRIMARY_TABLE_ID.

- 5 Field OBJ_NAME defines the name of the Docking Object being described. Field PRIMARY_TABLE_ID is used to identify the primary table associated with this Docking Object.

- 10 Table S_DOBJ_INST 63 describes whether a particular instance of a Docking Object, described by table S_DOBJ 61, is present on a particular node's database. Table S_DOBJ_INST 63 includes the fields NODE_ID, DOBJ_ID and PR_TBL_ROW_ID. Field NODE_ID points to a particular node table 65. Field DOBJ_ID points to the Docking Object to which the Docking Object instance applies. Field PR_TBL_ROW_ID is used to select a particular
15 row in the Primary Table of the Docking Object. This value identifies the Docking Object instance.

- Table S_REL_DOBJ 67 describes the related Docking Objects of a particular Docking Object, described by table S_DOBJ 61. Table
20 S_REL_DOBJ 67 includes the fields DOBJ_ID, REL_DOBJ_ID, and SQL_STATEMENT. Field DOBJ_ID identifies the Docking Object that owns a particular related Docking Object. Field REL_DOBJ_ID identifies the related Docking Object that is owned by the Docking Object identified by DOBJ_ID. Field SQL_STATEMENT is an SQL statement that may be executed to obtain
25 the Primary ID value of the related Docking Object.

- Table S_DOBJ_TBL 69 describes the member tables of a particular Docking Object, described by table S_DOBJ 61. Table S_DOBJ_TBL 69 includes the fields DOBJ_ID, TBL_ID, and VIS_EVENT_FLG. Field
30 DOBJ_ID identifies the Docking Object that contains the member table

described by the row. Field TBL_ID identifies the particular table in the database that is the member table described by the row. Field VIS_EVENT_FLG is a flag that indicates whether a change to this Docking Object can result in a visibility event. A value of "Y" indicates that a change
5 can result in a visibility event; a value of "N" indicates that it cannot.

Table S_DOBJ_VIS_RULE 71 contains the visibility rules associated with a particular Docking Object. S_DOBJ_VIS_RULE 71 contains the fields DOBJ_ID, RULE_SEQUENCE, RULE_TYPE, SQL_STATEMENT and
10 CHECK_DOBJ_ID. Field DOBJ_ID identifies the Docking Object with which a particular visibility rule is associated. Field RULE_SEQUENCE is a sequence number that indicates the sequence, relative to other visibility rules in table S_DOBJ_VIS_RULE 71, in which the particular visibility rule should be run. RULE_TYPE specifies whether the particular visibility rule is of type
15 "R," indicating an SQL visibility rule or of type "O," indicating a Docking Object visibility rule.

If RULE_TYPE is equal to "R," field CHECK_DOBJ_ID is not meaningful, and field SQL_STATEMENT contains an SQL statement that is
20 evaluated using the Primary ROW-ID of the primary table associated with this Docking Object and a particular Node 21. If the SQL statement returns any records, the Docking Object is deemed to be visible to the Node 21 for which visibility is being determined.

25 If RULE_TYPE is equal to "O," both field CHECK_DOBJ_ID and field SQL_STATEMENT are meaningful. Field CHECK_DOBJ_ID specifies a docking object whose visibility should be determined. If the specified docking object is deemed to be visible, then the docking object associated with the visibility rule is also visible. Field SQL_STATEMENT contains a SQL
30 statement that, when executed, returns the Row-ID of the docking object

identified by CHECK_DOBJ_ID that corresponds to the docking object instance associated with the visibility rule.

Table S_APP_TBL 73 is an Application Table that describes all the tables used in a particular application. It is pointed to by table S_DOBJ_TBL 69 for each member table in a docking object, and by table S_DOBJ for the primary table in a docking object. S_APP_TBL 73 points to table S_APP_COL 75, which is an Application Column Table that describes the columns of data in a particular application. S_APP_TBL 73 points to table S_APP_COL 75 directly through a primary key and indirectly through such means as a Foreign Key Column Table 81, User Key Column Table 83, and Column Group Table 85. The relationship of an Application Table, Application Column Table, Foreign Key Column Table, User Key Column Table and Column Group Table are well known in the art and are not further described.

Update Processing

Figure 3 depicts steps performed by an update manager 31 such as update manager 31-a, 31-b or 31-c in updating a database, such as a node database 23-a, 23-b or 23-c, responsive to user input. Execution of update manager 31 begins in step 101. In step 103, the update manager 31 accepts from the user input 33 in the form of a command requesting that the data in database 23 be altered. The request may be in the form of a request to delete a row of a table, to add a row to a table, or to change the value of a cell at a particular column of a particular row in a table. In step 105, using a well-known means, the update manager 31 applies the requested update to database 23. In step 107, the update manager 31 creates a log record describing the update and writes it to update log 35.

The contents of a log record describe the update made. Each log record indicates the node identifier of the node making the update, an identification of

the table being updated, and an identification of the type of update being made, i.e., an insertion of a new row, a deletion of an existing row, or an update to an existing row. For an insertion, the log record additionally includes an identifier of the row being inserted, including its primary key and the values
5 of the other columns in the row. For a deletion, the log record identifies the primary key of the row being deleted. For an update, the log record identifies the primary key of the row being updated, the column within the row being updated, the old value of the cell at the addressed row and column, and the new value of the cell.

10

After writing a log record in step 107, the update processor exits for this update. The foregoing description of the update processing preferably includes additional steps not material to the present invention, for example, to assure authorization of the user to make the update, to stage and commit the write to
15 the database to allow for rollback in the event of software or hardware failure, and the like. These steps are well-known in the art and are not described further.

An update manager 11 executing in central computer system 1 operates
20 in an analogous manner, except that it updates central database 3 and writes its log records to central update log 11.

Docking Processing

Figure 4 depicts steps performed by a Docking Manager 25 such as
25 Docking Manager 25-a, 25-b or 25-c to transmit and/or receive one or more transaction logs. Docking Manager 25 is invoked by the user of a remote node such as node 21-a, 21-b or 21-c, whereby the user requests that the node dock with central computer 1 to upload an update log such as update log 35-a to central computer 1, to download a partial log such as partial log 17-a, or both.
30 Execution of Docking Manager 25 begins in step 121. In step 123, Docking

Manager 25 connects with central computer 1 under the control of Central Docking Manager 5. This connection can be any connection that enables data exchange. It is anticipated that the most common form of a connection is a telephone line used in conjunction with a modem, but other forms of data connection, such as a Local Area Network or a TCP/IP connection may also be used. Step 125 checks to see whether the user has requested that node update log 35-a be uploaded to the Central Computer 1. If so, execution proceeds to step 127. If not, step 127 is skipped and control is given to step 129. In step 127, Docking Manager 25 uploads its update log to central computer 1. The upload may be accomplished with any known file transfer means, such as XMODEM, ZMODEM, KERMIT, FTP, ASCII transfer, or any other method of transmitting data. In step 129, Docking Manager 25 checks to see whether the user has requested that a partial log such as partial log 17-a be downloaded from Central Computer 1. If so, execution proceeds to step 131. If not, step 131 is skipped and control is given to step 133. In step 131, Docking Manager 25 downloads its partial log from central computer 1. The download may be accomplished with any known file transfer means, such as XMODEM, ZMODEM, KERMIT, FTP, ASCII transfer, or any other method of transmitting data. In step 133, having completed the requested data transfer, Docking Manager 25 exits.

Merge Processing

Merge processing is performed by a processor such as node merge processor 27-a, 27-b, or 27-c, or central merge processor 7. The merge process serves to update its associated database with a transaction that has been entered by a user of a computer remote from the computer where merge processing is being performed. Merge processing is analogous to update processing and is similar in form to update processing as previously disclosed with reference to figure 3, with three differences. First, the input to a merge processor is not an update entered directly by a user, but rather is a log file that is obtained from

a computer remote from the computer where the merge is executing. A second difference is that, as shown by in Figure 1, merge processing does not produce a log when performed at a node. The function of a log on a node is to record a transaction for propagation to Central Computer system 1 and thence to other nodes as required. A transaction that is the subject of a merge in a node has been communicated to Central Computer System 1, and there is no need to re-communicate it.

A third difference is that merge processing must be capable of detecting and resolving multiple conflicting transactions. For example, assume that a field contains the value "Keith Palmer." Assume further that a user at node 27-a enters a transaction to update that field to "Carl Lake," and a user at node 27-b enters a transaction to update the same field to "Greg Emerson." Without collision detection, data among various nodes may become corrupt. When the transaction for user 27-a is merged, the field is updated from "Keith Palmer" to "Carl Lake." Without collision handling, when the transaction for node 27-b is merged, the field would be updated to "Greg Emerson," and the central database would then be out of synch with the database of node 27-a. Furthermore, when merge processing is performed on each of nodes 27-a and 27-b, each node will update its database with the other's transactions, leaving at least one node out of synch with the other node and with central database.

Therefore, merge processing must also have a means of detecting collisions and correcting them. In the above example, a simple way to detect and correct a collision is to compare the value in the database to the value that the merge log reflects as being the previous value in the node database. If the two values do not match, Merge processor 7 may reject the transaction and generate a corrective transaction to be sent to the node from which the conflicting transaction originated. In the above example, when the transaction for node 27-b was presented to merge processor 7, merge processor 7 would

compare "Keith Palmer," the prior value of the field as recorded by node 27-b to "Carl Lake," the present value of the field as recorded in central database 3. Detecting the mismatch, merge processor 7 may then generate a transaction to change the value "Greg Emerson" to "Carl Lake," and write that transaction to update log 15. In a subsequent docking operation, that transaction would be routed back to node 27-b to bring its database 23-b in synch with the other databases.

The above is one example of a collision and a resulting corrective action. Other types of collisions include, for example, an update to a row that has previously been deleted, inserting a row that has previously been inserted, and the like. Merge processing must detect and correct each of these collisions. This may be performed using any of a number of well-known methods, and is not discussed further.

Figure 5 depicts the steps performed by merge processor such as central merge processor 7. Although it depicts merge processor 7 writing to central database 3 and to transaction log 15, it is equally representative of a node merge processor such as node merge processor 27-a, 27-b or 27-c updating a node database 23-a, 23-b or 23-c. Merge processing begins at step 141. In step 143, merge processor 7 finds the first unprocessed transaction on received log 19. In step 147, merge processor 7 selects a transaction from received log 19. In step 149, merge processor 149 attempts to update database 3 according to the transaction selected in step 147. In step 151, merge processor 7 determines whether the database update of step 149 failed due to a collision. If so, merge processor proceeds to step 153, which generates a corrective transaction. Following the generation of the corrective transaction, the merge processor returns to step 149 and again attempts to update database 3. If no collision was detected in step 151, execution proceeds to step 157. In step 157, merge processing checks to see if it is executing on central computer 1. If so,

step 155 is executed to journal the transaction to log 15. In any case, either if step 157 determines that the merge processing is being performed on a node or after step 155, execution proceeds to step 159. Step 159 checks to see if any transactions remain to be processed from log 19. If so, execution repeats from
5 step 147, where the next transaction is selected. If not, merge processing exits in step 161.

Log Management

Figure 6 depicts the steps to be performed by log manager 9 to prepare
10 a partial transaction log such as partial transaction log 17-a, 17-b, or 17-c. The procedure depicted in Figure 6 is executed for each node available to dock with central computer system 1. Log manager 9 begins execution in step 171. In step 173, Log Manager 9 finds the first unprocessed transaction for the node whose partial transaction log is being prepared. In step 175, log manager 9
15 selects a transaction for processing. In step 177, log manager 9 checks to see whether the selected transaction originated on the same node for which processing is being performed. If so, there is no need to route the transaction back to the node, and control proceeds to step 179. Step 179 checks to see whether there are any transactions remaining to be processed. If so, control is
20 given again to step 175. If not, control passes to step 189, which records the last transaction that was processed for this node, and then exits at step 191. If the transaction originates in other than the same node as the node for which processing is being performed, control is given to step 181. Step 181 calls a visibility calculator to determine whether the selected transaction is visible to
25 the node being processed. The Visibility calculator routine is described in detail further herein. In step 183, merge processor 9 checks to see whether the visibility calculator determined that the transaction is visible. If it is not visible, control is passed to step 179, which performs as disclosed above. If the transaction is visible, control is passed to step 185. Step 185 writes a
30 record for this transaction to the partial transaction log for the node being

processed, for example, partial transaction log 17-a for node 21-a. In step 187, the log manager 9 records the last transaction that was processed for this node, and then passes control to step 179, which determines whether to select additional transactions or exit, as disclosed above.

5

Visibility Calculation

Figure 7 depicts a flowchart describing the process a visibility calculator for calculating visibility for a docking object as invoked by step 181 of log manager 9. The visibility calculator is called with the node-id of the node for which visibility is being calculated, the docking object for which the visibility is being calculated, and the row-id of the docking object whose visibility id being calculated. The visibility calculator uses this information, in conjunction with information obtained from meta-data stored in the schema depicted in Figure 2, to determine whether a particular transaction that updates a particular row of a particular docking object is visible to a particular node.

The Visibility calculator begins execution at step 201. In step 203, the visibility calculator makes a default finding that the transaction is not visible. Therefore, unless the visibility calculator determines that a transaction is visible, it will exit with a finding of no visibility. In step 205, the visibility calculator selects the first visibility rule associated with the docking object. This is done by finding the table S_DOBJ_VIS_RULE 71 associated with the current Docking Object as pointed to by table S_DOBJ 61. In step 205, the visibility calculator selects the row of table S_DOBJ_VIS_RULE 71 with the lowest value for field RULE_SEQUENCE.

In step 207, the Visibility Calculator checks the field RULE_TYPE for a value of "R." The value of "R" indicates that the rule is a SQL visibility rule. If so, the Visibility Calculator proceeds to step 209. In step 209 the Visibility Calculator obtains a SQL statement from field SQL_STATEMENT

and executes it. An example of such an SQL statement might be:

```
5      SELECT 'X' FROM S_OPTY_EMP
      WHERE OPTY_ID = :PrimaryRowId
      AND EMP_ID = :NodeId;
```

This SQL statement causes a query to be made of application table S_OPTY_EMP. The query selects any records meeting two criteria. First, the records selected must have a field OPTY_ID, which is a row id or key, equal
10 to the Primary Row-ID of the Docking Object whose visibility is being determined. Second, the records selected must have a field EMP_ID, which may be for example, an identifier of a particular employee, equal to the NodeId of the node for whom visibility is being determined. In ordinary language, this SQL statement will return records only if a row is found in a table that matches
15 employees to opportunities, where the opportunity is equal to the one being updated, and the employee to whom the opportunity is assigned is the operator of the node.

This is a simplistic example, provided for maximum comprehension.

20 More complex SQL statements are possible. For example, the rule:

```
      SELECT 'X' FROM
      &Table_Owner.S_ACCT_POSTN ap
      &Table_Owner.S_EMP_POSTN ep
      WHERE ap.POSITION_ID = ep.POSITION_ID
25      AND ep.EMP_ID = :NodeId;
```

This rule queries the tables S_ACCT_POSTN (which relates a particular

account with a particular position in the organization that is responsible for the account) and S_EMP_POSTN (which relates what employee corresponds to a particular position). The condition "ap.POSITION_ID = ep.POSITION_ID" requires finding a row in the account-to-position table that has the same position
5 as a row in the employee-to-position table. The condition "ep.EMP_ID = :NodeId" further requires that the selected row in the employee-to-position table also have an Employee ID equal to the ID of the user of the Node for which visibility is being determined. In ordinary language, this condition allows visibility if the employee occupies the position that has responsibility for the
10 account in the docking object being updated.

There is no particular limit to the complexity of the conditions in the SQL statement used to evaluate visibility. Particular implementations of SQL may impose limitations, and resource considerations may make it desirable to
15 use less complex statements, but these limitations are not inherent in the invention.

Step 211 evaluates whether the execution of SQL_STATEMENT in step
209 returned any records. If records were returned, this indicates that the
20 Node for which visibility is being checked has visibility to the docking object being processed. Accordingly, if records are returned, the Visibility Calculator proceeds to step 213. In step 213, the transaction is marked visible. Because

no further rules need to be evaluated to determine visibility, the visibility calculator proceeds to step 228. Step 228 synchronizes the databases by determining whether the calculated visibility requires the insertion or deletion of a docking object into a particular node's partially replicated database. This
5 may occur, for example, if a node is determined to have visibility to a docking object due to a change to a related docking object. For example, an owner of a node may be assigned to a particular activity that is related to a particular sales opportunity. As a result, the node should be provided with a copy of the object representing the sales opportunity.

10

Figure 8 depicts the steps performed to synchronize a partially replicated database in response to a change in data visibility. Execution begins in step 241. In step 243, the Visibility Calculator references the visibility just calculated for a docking object. If the Docking Object is visible, execution
15 proceeds to step 245. Step 245 references the S_DOBJ_INST table, to verify that a row exists for the Docking Object for the current node. If a row exists, this indicates that the node in question already has a copy of the referenced Docking Object, and the routine proceeds to step 255, where it exits. If, however, no row exists for the Docking Object at the node being processes,
20 this indicates that the node in question does not have a copy of the Docking Object on its partially replicated database. The routine then proceeds to step 247, where a transaction is generated to direct the node to insert the Docking

Object into its partially replicated database.

If step 243 determines that the Docking Object is not visible, execution proceeds to step 249. Step 249 references the S_DOBJ_INST table, to verify
5 that no row exists for the Docking Object for the current node. If step 243 determines that no row exists in the S_DOBJ_INST table for the current docking object for the current row, this indicates that the node in question does not have a copy of the referenced Docking Object, and the routine proceeds to step 255, where it exits. If, however, a row exists for the Docking Object at
10 the node being processed, this indicates that the node in question does have a copy of the Docking Object on its partially replicated database. The routine then proceeds to step 251, where a transaction is generated to direct the node to delete the Docking Object from its partially replicated database.

15 Referring again to Figure 7, following the data synchronization routine of step 228, the Visibility Calculator proceeds to step 229, where it exits. Referring to Figure 6, as previously described, the resulting finding of visibility is available to be checked by the log manager in step 183 to determine to write the transaction.

20

Referring again to figure 7, if step 211 determines that no records were returned by the execution of the SQL statement in step 209, execution proceeds

with step 215. Step 215 checks to see whether there are any remaining visibility rules to be assessed. If not, the visibility calculator proceeds to step 228 to synchronize the database, and then to step 229, where it exits. In this case, the default mark of no visibility that was set in step 203 remains set.

- 5 This value will also be used by the log manager as shown in Figure 6, step 183, to determine not to write the transaction.

Referring again to Figure 7, if rules remain to be assessed, control proceeds to step 217, which selects the next rule to be processed. Control is
10 then given again to step 207 to begin processing the new rule.

The preceding text provided a description of the processing of SQL visibility rule; that is, visibility rules of type "R." If step 207 determines that the visibility rule is not of type "R," the visibility rule is of type "O." Type
15 "O" indicates a docking-object visibility rule. In such a case, the docking object being processed will be considered to be visible if it is related to a particular related docking object that is visible. If field RULE_TYPE is not equal to "R," then, execution proceeds to step 221. Step 221 determines the related Docking Object whose visibility must be determined to determine
20 whether the current docking object is visible. The related Docking Object identifier is obtained from field CHECK_DOBJ_ID in table S_DOBJ_VIS_RULE 71. In step 223, the Visibility Calculator determines

which row in the related Docking Object must be queried for visibility. In order to determine this, the Visibility Calculator obtains a predetermined SQL statement from the field SQL_STATEMENT and executes it. The SQL statement is a query that select one or more rows of the Docking Object that, for example, correspond to the docking object for which the Visibility Calculator was invoked.

For example, assume that it is desired to indicate that a record for a sales opportunity should be visible if the Node has visibility to any sales quote made for that sales opportunity. This may be accomplished using the following SQL statement:

```
SELECT "_ID" FROM
      &Table_Owner.S_DOC_QUOTE
      WHERE OPTY_ID=:Primary RowId
```

This SQL statement accesses a table S_DOC_QUOTE that contains all sales quotes. The WHERE clause specifies retrieval of all rows where the Opportunity ID of the row is equal to the Row-ID of the opportunity for which visibility is being calculated. The Visibility manager retrieves the specified Row-Ids, thereby identifying the rows of the S_DOC_QUOTE table whose visibility must be checked.

Having determined the a related docking object and the row-ID of that related docking object upon whose visibility the visibility of the current docking

object depends, the Visibility Calculator proceeds to step 225. In step 225, the Visibility Calculator recursively invokes itself to determine visibility of the related docking object. The recursively invoked Visibility Calculator operates in the same manner as the Visibility Calculator as called from the Log Manager

5 9, including the capability to further recursively invoke itself. When the recursive call concludes, it returns a visibility indicator for the related Docking Object, and control proceeds to step 227. In step 227, the Visibility calculator determines whether the related Docking Object was determined to have been visible. If so, the Visibility Calculator proceeds to step 213 to mark the

10 originally current Docking Object as visible, and then to step 228 to synchronize the database and then to step 229 to exit. If the related Docking Object was not determined to be visible, control proceeds to step 215 to determine whether additional visibility rules remain to be assessed.

The Visibility Calculator, in conjunction with the Log Manager is

15 therefore able to determine what subset of update transaction data is required to be routed to any particular node. This operation serves to reduce the transmission of unneeded data from the Central Computer 1 to the various nodes such as nodes 21-a, 21-b and 21-c that utilize partially replicated databases, and to reduce the system resources such as disk space needed to

20 store, and the CPU time needed to process, what would otherwise be required to maintain a fully replicated database on each remote node.

The operation of the log manager 9 in conjunction with the Visibility Calculator herein described will be apparent from reference to the description and to the drawings. However, as a further aid in the description of these facilities, a pseudocode representation of these facilities is hereto attached as
5 an Appendix.

Internet Session Manager

The following terms will be used in describing the Internet Session Manager:

- 10 • SISM: Siebel Internet Session Manager (also referred to as Internet Session Manager).
- SISMGR: The application server which implements SISM (the program itself).
- SIS API: The client-side application programming interface for
15 communicating with the SISM server.

The Internet Session Manager provides database functionality and access to mobile clients that are connecting via a wide area network such as the Internet, a corporate intranet, or a commercial online provider. The
20 functionality and access thus provided to mobile clients is comparable to that of clients connected to the central database via a local area network. The Internet Session Manager also provides security for data that flows between the

client and the database, as will be described herein.

Overview of Internet Session Manager

SISMGR is a new application server process which currently runs on
5 Windows NT, but may be ported to later versions of Windows NT and to
UNIX. One or more instances of SISMGR run continuously on a machine to
accept connections from remote clients. Each client connection "into" SISMGR
corresponds to a database connection "from" SISMGR, which essentially acts
as the connector between the client and the database. Thus, a machine running
10 SISMGR is configured as a networked proxy server. The overall structure of
system using the Internet Session Manager is shown in Figure 9. SISMGR has
several advantages over a direct SQL connection such as Oracle's SQL*Net):

1. Security for data which flows between the client and server.
2. True access control for client authentication.
- 15 3. Central monitoring and control of remote access.

Data Security

SISMGR addresses several methods of intercepting the data transferred
between the client and SISMGR:

- 20 4. No one besides the two parties (mobile client and SISMGR) should be
able to read the contents of the transmissions.
5. No one should be able to alter the data during transmission and have the

results accepted.

6. No one should be able to replay parts of a previous transmission and have the results accepted.

5 These three types of attacks are primarily addressed through encryption.

Encryption techniques fall into two basic categories: *public key* and *private key* cryptography. Private key cryptography uses a key which must be kept secret (hence the "private"), and is very secure as long as the key is not compromised. Encryption and decryption are rapid, and private key encryption
10 is used to transmit the bulk of the data in secure systems. Public key cryptography techniques provide encryption and decryption through a pair of keys, designated as the "private key" and the "public key" (although they are interchangeable). The power of this technique is that one key can be given to the world and one kept secret. The "trick" of public key cryptography is that
15 data encrypted with the public key can only be decrypted with the private key and vice versa. Thus a message can be encrypted (with the public key) which can only be decrypted by the intended recipient and a message can be verified as having been sent by a person if it can be decrypted with that person's public key. Further discussion may be found in "PGP: Pretty Good Privacy" by
20 Simson Garfinkel, and U.S. Patents No. 4,405,829, 4,218,582, 4,424,414, and 4,200,770, the disclosures of which are hereby incorporated by reference.

SISMGR uses a session key and sequencing to provide a secure communication channel. The first packet sent by the client contains the client's public key, which can be verified by the server against an internal database or trusted certification authority (CA). The server reply contains a session
5 identifier (unique for each communication session) and a session private key (randomly generated by the server) and will be encrypted with the client's public key. All further messages in both directions are encrypted with the session key using private key cryptography. All messages will contain the non-repeating session identifier allocated by the server, which prevents messages
10 from being replayed in another session. Each message in a session has a monotonically increasing identifier to prevent messages from being replayed in or removed from the session.

Access Control

15 In addition to preventing communications between legitimate users from being intercepted, it is necessary to prevent illegitimate users from gaining access to SISMGR. This is especially important in implementations where all Visibility Rule checking is performed by the client and access to SISMGR provides full access to the corporate database, but is still of importance in
20 implementations where Visibility Rule checking is performed by the Central Computer 1.

Access control prevents both illegitimate use of the Siebel client application as well as special programs which pretend to be clients. As in all security systems, the weak link is keeping the private key from getting out. Thus, the private key is stored in encrypted form on the user's laptop, and the user must enter a password (or phrase) to unlock his/her private key. Otherwise, anyone with access to the laptop would have access to the corporate database. The user must still enter his/her database name and password to actually log in.

One skilled in the art will recognize that hardware-based authentication schemes, such as those which use PCMCIA or flash memory cards to hold the keys and possibly perform the encryption, may be used without departing from the spirit or scope of the present invention.

15 *Monitoring and Control*

The SISMGR keeps a list of hosts which are allowed to connect. Connections from hosts or network addresses which are not listed will be rejected. This helps avoid attacks from systems which should not have access. An access control list is of particular utility on an intranet where all points of access are known in advance. Further details of the access control list will be described below.

In an alternative embodiment of the invention, SISMGR may include security auditing tools which check for patterns of behavior which might indicate attacks on the system, such as logins outside typical working hours, many failed login attempts, or sessions with atypical query patterns. Further, there could be a mechanism for evaluating the load and response time for clients, to enable system administrators to ensure good throughput at peak times.

Architecture


The functionality of the Internet Session Manager is implemented in a common DLL (ssasis30.dll) and a number of connector DLLs (such as ssx30o72.dll for Oracle). A simple executable wraps this functionality to produce a server. SISMGR is a main routine which links in and initializes the SISM DLL and displays errors on the console. In another embodiment, this functionality may be implemented as a Windows NT service which has no user interface at all.

Also part of the SISM functionality is the "SIS network API," also a DLL (ssasna30.dll) which encodes and decodes the messages which flow back and forth between the mobile client and the SISMGR. The actual messages, their parameters, and the actual format of the data are all defined in the SIS network API.

As will be described below, the present embodiment of the invention defines a set of messages which allow the client to implement a connector which executes its

SQL operations using the SIS network API to talk to the SISMGR. However, the architecture may be extended to provide further functionality for mobile clients. For example, an "object manager" server may be used to implement a "thin client" interface and even alternate clients.

5



This implementation essentially moves the actual database connection (the code which talks to the database vendor APIs) out of the client process and into the SISMGR process. This permits the imposition of additional restrictions and provides complete central control over the data which passes across the connection between the client and the database server.

SISMGR provides a session-oriented protocol as part of the security setup. The first message sent by a client must be a HELLO message, which establishes the client authentication and creates a session ID which is used by all remaining requests in this session. The details of the HELLO message (and all other messages) can be found in the `sisnapi.h` header file below.

SISMGR supports direct SQL access messages which provide direct support for the client connector operations. The messages in this group all translate directly into database operations, although a single message may logically correspond to multiple database API calls. Before any other SQL operations can be performed, the user must log into the SQL database. This is done with the LOGIN message, which takes database connection information. The LOGOUT message terminates a login session. Only a single login session may be active at once for a client.

Actual SQL statements are always executed against a "statement context."
This statement context corresponds to an ODBC HSTMT or an Oracle cursor.
(ODBC, or Open DataBase Connectivity, is a specification for an API that enables
applications to access multiple database management systems using SQL. HSTMT
5 a statement handle.) Before an actual SQL statement can be executed, a statement
must be allocated with the ALLOCSTMT message (which returns a HSTMT).
Statements can be freed with the FREESTMT message or simply re-used for another
SQL statement. These statement IDs must be specified for executing statements and
fetching rows.

10

The EXEC message is used to execute a SQL statement (usually a DML, or
Data Manipulation Language, statement) one or more times with different bind
variables. Because no output is returned from the statement, this cannot be used for
queries. However, the ability to re-execute the statement for each set of bind
15 variables is very useful for reducing round trips for repetitive operations. The
QUERY and FETCH messages provide query support. QUERY is sent to begin a
new query statement and return the first batch of results. FETCH is used to get the
next batch of results. The size of the result set can be specified in QUERY to limit
the number of rows fetched by each request. The COMMIT and ROLLBACK
20 messages are provided to commit and roll back database transaction operations. The
EXEC message can also request that a commit be performed if the operation is
successful without requiring an explicit COMMIT message be sent.

To support literature and correspondence (fulfillment), SISMGR provides
25 simple file services. Basic file-oriented operations can be performed for files by
specifying file names relative to the root directory configured into the SISMGR.
Information about files can be obtained by using the FILEINFO message, which
returns standard file information but does not actually open the file. To open a file
and allocate a "file ID" which is used by the other operations on this section, the
30 FILEOPEN message is used. FILEOPEN returns roughly the same information as
returned by FILEINFO, plus the file ID. Files are closed after use with the
FILECLOSE message. Files are read from and written to using the FILEREAD and

FILEWRITE messages. These operate very similarly to the standard file operations and their use is well known in the art. In the present embodiment, one restriction is that a file may only be opened for reading or writing, not both. Opening a file for write always truncates the existing contents (if any) and creates the file if it does not exist.

Also provided are high-speed (tuned) connectors for the different databases as part of SISM, in the same way that these connectors are provided in the client. The SISM connectors are simpler than the client connectors because they do not need to know how to format statements. They need only execute them, collect results, and handle errors.

The abstract class C~~SS~~SISDbConn (declared in `siscon.h`) defines the template which each connector function must follow, and defines the interface which database-specific connectors must implement to support the SISM client database operations. The templates correspond roughly to the ANSI cursor model, but with some compression to reduce round trips. The data passed to the functions is taken directly out of the messages with little translation by SISMGR, but the connectors are not themselves dependent on the message format.

20

An instance of a sub-class of C~~SS~~SISDbConn is created for each client connection when it attempts to log in. The client passes information about which database it wants to log into and the user name and password of the user to log in as. SISMGR maps the requested "database" into the database-specific information as it validates the request.

25

Each client request is translated into one or more calls to this API, which is structured slightly differently from the messages for simplicity. The main goal of the message API is to minimize round trips. The main goal of the connector API is to be easy to implement.

30

SISMGR is a multi-threaded program, and each client instance runs in its own

thread. As is the case everywhere else in SISMGR, the connector implementations must be careful to be multi-threading-safe. If necessary, synchronization objects implemented in SISMGR can be used, but this will probably reduce throughput when many clients are connected.

5

Further details of C\$SSISDbConn, related code, and the functionality may be seen in the following sismcon.h code:

```

10  // FILE:      sismcon.h
    // DESCRIPTION
    //
    //   Generic DB-specific connector class for Internet Session Manager.

15  #ifndef __SISMCON_H__
    #define __SISMCON_H__

    #ifndef __SISNAPI_H__
        #include "sisnapi.h"
    #endif

20  class C$SSISDbStmt;

    #define MAX_NUM_STMTS      256

25  class SS_EXPORT_CLASS C$SSISDbConn : public C$SSISObject
    {
    protected:
        BOOL      m_bConnected;
        C$SSISDbStmt* m_dbStmtArray [MAX_NUM_STMTS];

30  public:
        //
        //   When the client sends a LOGIN request, the Login method is
35  called
        //   this should perform any database-specific login operations and
        //   return an appropriate error status. Logoff performs the reverse
        //   operations and frees all resources associated with the previous
        //   login.
40  //
        //   Note that since there is an instance of a concrete sub-class of
        //   C$SSISDbConn for each client, the database login is one-to-one
        //   with the instance of C$SSISDbConn. Thus both session and login
        //   information should be stored in members of the sub-class.
45  //
        virtual BOOL      DoLogon (const char* pInstance,
                                   const char* pUserName,
                                   const char* pPassword) = 0;

50  virtual BOOL      Logon (const char* pInstance,
                            const char* pUserName,
                            const char* pPassword);

        virtual BOOL      DoLogoff (void) = 0;
        virtual BOOL      Logoff (void);

55  //
        //   The statement executing APIs parallel the SISM request messages
        //   closely. The Execute call prepares a SQL statement and then

```

```

// executes it a specified number of times, each time with a
// different set of bind variables. If no bind variables are
// given, the statement should be executed once.
5 //
// No output values are expected from the Execute statement and
// no persistent cursor information should be maintained. The
// number of rows affected (a sum of all executions) should be
// returned to the caller through the last parameter.
//
10 virtual BOOL      Execute (SISStmtId      stmtId,
                           CSSSISColRowData* pIBinds,
                           int&              numAffected);

15 virtual BOOL      ExecuteAndFetch (SISStmtId      stmtId,
                                   CSSSISColRowData* pIBinds,
                                   CSSSISColRowData* pFetched,
                                   int&              numFetched);

20 //
// For statements which return output (SELECT statements), we use
// the ANSI model of preparing the statement and executing it and
// then fetching results. The prepare must open a cursor which
// is returned to the caller for future calls to Fetch and finally
25 // a call to Close.
//
// The fetch function takes an argument of the number of rows to
// fetch (at most) and an indicator for whether this drains the
// cursor. This information is returned to the client.
//
30 virtual BOOL      Prepare (SISStmtId      stmtId,
                           const char*      pStmt,
                           CSSSISColRowData* pOBinds,
                           int              nMaxFetch);

35 virtual BOOL      Fetch (SISStmtId      stmtId,
                           CSSSISColRowData* pFetched,
                           int&              numFetched,
                           BOOL&            bEOF);

40 virtual BOOL      AllocStmt (SISStmtId&   stmtId) = 0;
virtual void        DoFreeStmt (CSSSISDbStmt* pDbStmt) = 0;
virtual BOOL        FreeStmt (SISStmtId     stmtId);

45 //
// For maintaining the transaction state, the Commit and Rollback
// methods perform the DB operations.
//
50 virtual BOOL      Commit (void) = 0;
virtual BOOL        Rollback (void) = 0;

virtual void        Release (void) = 0;

55 protected:
                                CSSSISDbConn (void);
virtual            ~CSSSISDbConn (void);

60 private:
    BOOL GetStmt (SISStmtId      stmtId,
                  CSSSISDbStmt* pDbStmt);

65 };

```

```

class SS_EXPORT_CLASS CSSSISDbStmt: public C$$$ISObject
{
protected:
5   C$$$ISDbConn*      m_pDbConn;
   C$$$ISColRowData*  m_pOutputBindColData;

public:
10   C$$$ISDbStmt (C$$$ISDbConn* pDbConn);
   virtual ~C$$$ISDbStmt (void);

   virtual BOOL      DoExecuteAndFetch (C$$$ISColRowData* pIBinds,
15   int& numFetched) = 0;
   virtual BOOL      DoFetch            (int& numFetched,
   BOOL& bEOF) = 0;
   virtual BOOL      DoPrepare          (const char* pStmt,
   C$$$ISColRowData* pOBinds,
20   int nMaxFetch) = 0;
   virtual BOOL      Execute            (C$$$ISColRowData* pIBinds,
   int& numAffected) = 0;
   BOOL      ExecuteAndFetch (C$$$ISColRowData* pIBinds,
   C$$$ISColRowData*& pFetched,
25   int& numFetched);
   BOOL      Fetch                (C$$$ISColRowData*& pFetched,
   int& numFetched,
   BOOL& bEOF);
   BOOL      Prepare              (const char* pStmt,
   C$$$ISColRowData* pOBinds,
30   int nMaxFetch);
};

//
// Each DLL which implements a specific database connector must have
// an exported function named Create$$$ISDbConn which takes as its
35 single
// argument a reference to a C$$$ISDbConn* into which it will return
// the constructed instance of a sub-class of C$$$ISDbConn.
//
// The typedef here gives the format of the function which must be
40 // implemented by the connector. All other methods implemented by
// the connector are done as methods on the constructed C$$$ISDbConn
// instance.
//
45 typedef BOOL (*$$$IS_CREATE_FN) (C$$$ISDbConn*&);

#endif // !__$$$ISCON_H__

```

50 One important data structure defined in the SIS network API is C\$\$\$ISColRowData (declared in `sisnapi.h`) which is the buffer which holds rows and columns of data input bind variables to SQL statements and rows of data output by a query. For efficiency, these data structures are passed directly back and forth between the connector methods and the messages sent and received by SISMGR.

55 This avoids excess data copying. Further details of C\$\$\$ISColRowData and other information will be apparent from inspection of `sisnapi.h`, which defines the structure of messages that flow back and forth between the client and the SISM

server. The protocol is structured as request/reply (remote procedure call). The client may implement a non-blocking model to avoid "locking up" while waiting for the reply, but a simple client implementation could treat these as regular procedure calls, not requiring a major change in architecture.

5 The messages have a simple format common format made up of a constant header and an optional list of attributes. The attributes are message specific and list the elements of the message which are required.

10 At the lowest level, all data are represented as either integers or strings. All integer data are in "network byte order" to avoid interpretation problems on heterogeneous architectures. All string data are an array of bytes (which may include the NUL byte), although usually these will be true C strings which are NUL terminated. Strings are always byte-oriented since the data is not byte swapped for transmission.

15 The length of the header and the length of the body are both explicitly given. This allows for advances in the protocol to be made without breaking older implementations (unless desired) as long as new fields are appended to the existing messages.

20 The request, and associated replies contain the header plus additional information. All negative acknowledgement (NAK) messages have the same format: a generic error code (defined in `sisnapi.h`) and optional detail information.

25 When the message is actually transmitted, the entire message (both the header and body) are compressed and encrypted (except for the HELLO request and NAK messages and the NOACCEPT notify message). Because of this, the length of the entire message (after compression and encryption) are sent as a `SISUInt32` (in the clear) immediately before the message data and are not included in the header length.

30 In order to make it easier for the client and server to construct and parse the messages sent back and forth, a set of classes are given to pack and unpack the

messages and perform the compression and encryption as necessary. The abstract superclass of these is C\$\$\$SISMessage, which has four sub-classes (one for each message type). Following is the sisnapi.h file, showing the details of the C\$\$\$SISMessage and other structures.

```

5  // FILE:      sisnapi.h
   // DESCRIPTION
   //
   //   Public header file for the Siebel Internet Session message API.
10 #ifndef __SISNAPI_H__
   #define __SISNAPI_H__

   class C$$$SISColRowData;
15   enum SISMessageType
   {
       SIS_MSG_UNKNOWN    = 0,
       SIS_MSG_REQUEST    = 1,    // client request message
20   SIS_MSG_ACK          = 2,    // server success message
       SIS_MSG_NAK        = 3,    // server failure message
       SIS_MSG_NOTIFY     = 4     // server status message
   };

25   enum SISRequestCode;

   class SS_EXPORT_CLASS C$$$SISMessage
   {
   public:
30   SISMessageType        GetType () const { return m_type; }

       //
       //   Methods for constructing a message.
       //   Firstly, the message type is given, which allows the header to
35   be
       //   constructed. Then, data is added according to the actual
       //   contents
       //   of the message.
       //
40   protected:
       BOOL                Add_int (unsigned int value);

       BOOL                Add_str (const char* pStr);

```

```

        BOOL                Add_str (const char*  pBuf,
                                   int             nLen);

        BOOL                Add_crd (CSSSISColRowData* pCRD);
5
    public:
        virtual BOOL        Encode (SISMessageType);

10        BOOL              Complete ();

        unsigned char*      GetData () { return m_pDataBuf; }
        int                 GetLength () const { return (m_pDataPtr -
15        m_pDataBuf); }

```

The sisnapi.h file also defines methods for unpacking a message. The message data is given, from which the message contents will be extracted. The CopyData method makes a copy of the given buffer while TakeData just keeps a pointer to the passed buffer and will free it when the message is deleted.

```

    public:
        BOOL                CopyData (unsigned char* pData,
25        int                 nLen);

        BOOL                TakeData (unsigned char* pData,
                                   int                 nLen);

        void                DropData ();

30        virtual BOOL      Decode ();

    public:
        CSSSISMessage ();
        ~CSSSISMessage ();
35        virtual

        void                Release ();

40    protected:
        BOOL                Get_int (unsigned int*  pBuf);

```

```

        BOOL                Get_str (char**      ppData,
                                   int*         pLen);

        BOOL                Get_crd (CSSSISColRowData** ppCRD);
5
    private:
        SISMessageType       m_type;
        unsigned char*       m_pDataBuf;
10        unsigned char*     m_pDataPtr;
        int                 m_nBufSize;
        int                 m_nHeaderLen;
        unsigned char*      m_pBodyLen;

15        BOOL                Extend (int nNeedBytes = -1);
    };

```

Further provided are functions to connect to the SISM server (or any TCP/IP server). These functions do not transmit or receive any data. They simply construct and tear down the TCP/IP connection. SISConnectServer returns a socket identifier which is passed to SISSendMessage and SISReadMessage to actually transmit and receive message data. Only a single server session created by SISConnectServer may be outstanding at any one time. This is because there is only a single buffer used by SISReadMessage for collecting the input as it arrives on the socket.

The function SISIsConnBroken may be called to return the status of the connection. If it returns TRUE, no more messages can be read or sent over this connection, and a new connection must be established (or the client should terminate).

30

```

//
extern int SS_EXPORT_API SISConnectServer (const char* pHostName,
                                           int         nInetPort);
35
extern void SS_EXPORT_API SISDisconnect (int nSocketId);

extern BOOL SS_EXPORT_API SISIsConnBroken (int nSocketId);

```

Security is also provided for SISMGR in the `sisnapi.h` file. The client transmits its public key using `SISCryptGetPublicKey()` in the HELLO message, in the clear. The server generates a session key and exports it, encrypted with the client's public key using `SISCryptGenSessionKey()`, and sends it back in the HELLO ACK, also in the clear. The HELLO and the HELLO ACK are the only messages in the clear. All other messages are encrypted with `SISCryptEncrypt()` and decrypted with `SISCryptDecrypt()`, using the session key. `SISCryptInitialize()` is called by the client before the HELLO message and by the server before the HELLO ACK to initialize the cryptography session. If required, a key container is created and a public key generated.

`SISCryptGetPublicKey()` is called by the client to get the public key prior to encoding the HELLO message.

`SISCryptImportKey()` is called by the client to import the session key on the HELLO ACK and by the server to import the client's public key in the HELLO message.

`SISCryptValidatePublicKey()` is called by the server on receipt of the HELLO message to validate and import the client's public key.

`SISCryptGenSessionKey()` is called by the server prior to sending the HELLO ACK to generate a session key, encrypted by the client's public key.

`SISCryptEncrypt()` and `SISCryptDecrypt()` work with the `CompressBuf` structure.

```
//
typedef unsigned long SISCryptProvider;    // HCRYPTPROV
30 typedef unsigned long SISCryptKey;      // HCRYPTKEY

extern BOOL SS_EXPORT_API SISCryptInitialize(
                                SISCryptProvider* phProv);    //
out
```

```

extern BOOL SS_EXPORT_API SISCryptGetPublicKey(
    SISCryptProvider hProv,           //
    in
    SISCryptKey*     phPublicKey,     //
5  out
    unsigned char**  ppszPublicKeyBlob, //
    out
    DWORD*           pdwPublicKeyLength); //
10 out
extern BOOL SS_EXPORT_API SISCryptImportKey(
    SISCryptProvider hProv,           //
    in
    SISCryptKey*     phKey,           //
15 out
    unsigned char*    pszKeyBlob,      //
    in
    DWORD             dwKeyLength);    //
20 in
extern BOOL SS_EXPORT_API SISCryptValidatePublicKey(
    SISCryptProvider hProv,           //
    in
    SISCryptKey*     phPublicKey,      //
25 out
    unsigned char*    pszPublicKeyBlob, //
    in
    DWORD             dwPublicKeyLength); //
30 in
extern BOOL SS_EXPORT_API SISCryptGenSessionKey(
    SISCryptProvider hProv,           //
    in
    SISCryptKey       hPublicKey,      //
35 in
    SISCryptKey*      phSessionKey,    //
    out
    unsigned char**    ppszSessionKeyBlob, //
    out
40 out
    DWORD*             pdwSessionKeyLength); //
    out
extern void SS_EXPORT_API SISCryptDestroyKey(
    SISCryptKey&      hKey);           //
45 in/out

```

```
extern void SS_EXPORT_API SISCryptRelease(  
                                     SISCryptProvider hProv );           //  
in
```

5

When a message comes in on a socket (bytes are ready to read), the function `SISReadMessage` is called. This function will block until the entire message has been read or until the connection is closed. The sub-class of `CSSSISMessage` appropriate to the received message is returned. The session key to be used to decrypt the message is passed in. If a message is in the clear, this is 0 (**HELLO**, **HELLO ACK**). If more bytes are read on the socket than are needed to complete the message, they will be saved in a buffer private to `SISReadMessage`. Because of this, this function is not thread safe and cannot be used by the `SISMGR` itself.

15

`SISDecodeMessage` takes the buffer given (which is already assumed to be a complete message) and decodes it as an actual message. This is the form used by the `SISMGR`. `SISReadMessage` reads the necessary data, then calls `SISDecodeMessage` to build the message structure. The buffer is kept by the message which is constructed and will be freed when the message is deleted. The caller must not modify or free the passed data because it is used as the actual message buffer. The session key to be used for decryption is passed in.

20

`SISCheckMessages` returns the number of messages which can be gotten using `SISReadMessage` without blocking. `SISCheckMessages` may actually read data, but will not block. If the connection is broken, `SISCheckMessages` returns -1.

25

When a message is to be sent, the appropriate concrete sub-class of `CSSSISMessage` is created and the data is encoded, using the virtual `Encode` method. Then the message is written onto the socket in one piece. As soon as this function returns, the message can be deleted. The session key to be used for message encryption is passed in. If this is 0, the message is sent in the clear. This happens only in **HELLO** and **HELLO ACK** msgs and also the server's `NoAccept()` response:

30

```

//
extern BOOL SS_EXPORT_API SISReadMessage (int          nSocketId,
                                           CSSSISMessage* pReturn,
                                           SISCryptKey  hSessionKey);
5
extern BOOL SS_EXPORT_API SISDecodeMessage (unsigned char* pBuf,
                                           int          nLen,
                                           CSSSISMessage* pReturn,
                                           SISCryptKey  hSessionKey);
10
extern int  SS_EXPORT_API SISCheckMessages (int nSocketId);

//
//
15
extern BOOL SS_EXPORT_API SISSendMessage (int          nSocketId,
                                           CSSSISMessage* pMsg,
                                           SISCryptKey  hSessionKey);

```

20

Each request has a unique code (SIS_REQ_*) which identifies the request (operation) from the client. Success and failure (ACK and NAK) reply messages will be returned for each request message. A simple success message might contain no special information, but a more complex one would contain results of the operation.

25

The format of all NAK messages is the same. When a request is constructed, the next "request sequence number" is assigned to it for sequence verification. This is done automatically when the request is created. The request ID is used to connect the ACK or NAK reply to its original request. This mechanism provides for allowing asynchronous requests.

30

```

//
typedef unsigned long  SISSessionId;
typedef unsigned long  SISSequence;
35
enum SISRequestCode
{
    SIS_REQ_HELLO      = 101,    // establish connection with server
    SIS_REQ_LOGON      = 102,    // attempt to log the user in to the
40    corporate DB
    SIS_REQ_LOGOFF     = 103,    // close DB connection for current user
    SIS_REQ_ALLOCSTMT = 201,    // allocate a statement for executing
45    DML or query
    SIS_REQ_EXEC       = 202,    // execute a DML statement--no query
    results
    SIS_REQ_QUERY      = 203,    // execute a query and return first
    rows
50    SIS_REQ_FETCH      = 204,    // fetch more results on the statement
    SIS_REQ_FREESTMT   = 205,    // free a statement opened by
    SIS_REQ_ALLOCSTMT
    SIS_REQ_COMMIT     = 206,    // commit operations on a session
55    SIS_REQ_ROLLBACK   = 207,    // rollback operations on a session

```



```

        SIS_REQ_FILEINFO = 208,    // return information about a file on
the server
        SIS_REQ_FILEOPEN = 209,    // open a file on the file server
        SIS_REQ_FILEREAD = 210,    // read a chunk of data from the given
5  file
        SIS_REQ_FILEWRITE = 211,   // write a chunk of data to the given
file
        SIS_REQ_FILECLOSE = 212,   // close a file opened with OPENFILE
};
10
class SS_EXPORT_CLASS CSSSISRequestMsg : public CSSSISMessage //
ABSTRACT
{
15  public:
        SISRequestCode      GetRequest () const { return m_request; }
        SISSessionId        GetSessionId () const { return m_sessionId; }
        SISSequence         GetSequence () const { return m_sequence; }

        virtual BOOL        Encode (SISRequestCode,
20                               SISSessionId,
                               SISSequence&);

        virtual BOOL        Decode ();

25  virtual                  CSSSISRequestMsg ();
        virtual              ~CSSSISRequestMsg ();

30  private:
        SISRequestCode      m_request;
        SISSessionId        m_sessionId;
        SISSequence         m_sequence;
};
35  //

```

“Positive acknowledgement” messages may contain additional information specific to the message to which they reply to. A generic ACK merely indicates that the request was processed successfully. Most of the actual ACKs add information to the basic ACK structure. Thus, the CSSSISACKMsg class is subclassed in parallel to the CSSSISRequestMsg class. Currently these are one-to-one.

All “negative acknowledgement” messages are the same format for simplicity. A NAK code is returned plus optionally some detail as a text message for the user. The NAK code allows the client to perform recovery as appropriate and the detail message may be displayed to the user if desired (probably most interesting for debugging). A single NAK message has an array of error indications within it. This represents the heirarchy of errors which are recognized by the software from general to specific. Often, there will only be a single error, but typically there will be errors from different levels. When constructing a NAK

message, the error information passed in is copied into the NAK struct and deleted when the NAK message is deleted. When unpacking a NAK message, the CSSSISErrorInfo structs should be read, but not modified or pointers kept because they will be deleted with the containing message.

```

5
//
class SS_EXPORT_CLASS CSSSISACKMsg : public CSSSISMessage //
ABSTRACT
{
10 public:
    SISSequence      GetRequestId () const { return m_requestId; }
    SISRequestCode    GetRequestType () const { return m_reqType; }

15     virtual BOOL    Encode (SISSequence      requestId,
                             SISRequestCode reqType);

    virtual BOOL      Decode ();

20     virtual          CSSSISACKMsg ();
                    ~CSSSISACKMsg ();

    private:
        SISSequence    m_requestId;
25         SISRequestCode m_reqType;
    };

//
//
30 enum SISErrorType
{
    SIS_NAK_INTERNAL    = 0,
    SIS_NAK_SISM         = 1,      // SISM errors (one of SIS_ERR_* below)
35     SIS_NAK_DENATIVE  = 2,      // database vendor specific error
    SIS_NAK_DBGENERIC    = 3      // database non-vendor specific error
                                // (one of SIS_DBGEN_ERR_* below)
};

40 enum SISErrorCode
{
    SIS_ERR_PARAMS      = 100,    // generic error in message params
    SIS_ERR_VERSION     = 101,    // protocol version not supported
    SIS_ERR_KEY         = 102,    // invalid client public key
45     SIS_ERR_AUTH      = 103,    // invalid authentication info.
    SIS_ERR_CONNECT     = 104,    // invalid connect string (DB instance)
    SIS_ERR_LIMIT       = 105,    // resource allocation not allowed
    SIS_ERR_REQUEST     = 106,    // unexpected request type
    SIS_ERR_STMTID      = 107,    // invalid statement ID given
50     SIS_ERR_ACCESS    = 108,    // access to system resource denied
    SIS_ERR_FILESYS     = 109,    // error opening/reading/writing file
    SIS_ERR_FILEID      = 110,    // invalid file ID given
    SIS_ERR_FILECAT     = 111,    // invalid file category used
    SIS_ERR_CRYPTINIT   = 112,    // error initializing Crypto API
55     SIS_ERR_SESSKEY   = 113    // error generating session key
};

enum SISDbGenErrorCode
60 {
    SIS_DBGEN_ERR_WRITECONFLICT = 101,

```

```

        SIS_DBGEN_ERR_DUPCONFLICT    = 102,
        SIS_DBGEN_ERR_INVALLOGON     = 103
    };

5
    class SS_EXPORT_CLASS CSSSISErrorInfo
    {
    public:
        SISErrorType      m_errType;
10        unsigned int     m_errCode;
        char*             m_pMessage;

        BOOL              SetErrorInfo (SISErrorType type,
15                                unsigned int code,
                                const char* pMessage);

                                CopyErrorInfo (const CSSSISErrorInfo* pCopy);

        void              Release ();

20                                CSSSISErrorInfo ();
                                ~CSSSISErrorInfo ();
    };

25    class SS_EXPORT_CLASS CSSSISNAKMsg : public CSSSISMessage // FINAL
    {
    public:
        SISSequence       GetRequestId () const { return m_requestId; }
30        int              GetErrorCount () const { return m_nErrInfos; }
        CSSSISErrorInfo*  GetErrorInfo (int nIndex);

        virtual BOOL      Encode (SISSequence request,
                                CSSSISErrorInfo** ppErrs,
                                int nErr);
35        virtual BOOL      Encode (SISSequence,
                                SISErrorType type,
                                unsigned int code,
                                const char* pMessage);

40        virtual BOOL      Decode ();

        static BOOL        NewNAK (CSSSISNAKMsg*& pReturn);

45        CSSSISNAKMsg ();
        virtual            ~CSSSISNAKMsg ();

50    private:
        SISSequence       m_requestId;
        CSSSISErrorInfo** m_ppErrInfos;
        int               m_nErrInfos;
55    };

```

In addition to the RPC protocol initiated by the client, the server can send messages to the client (SIS_NTY_*). These messages are not acknowledged by the client, and the server does not wait for a response. The client can choose to

60 handle or not handle these notifications (for example, notify the user).

The HELLO message contains information from the client which allows the server to begin communicating with it. The major piece of client information is the public key the client is using. The server will encrypt the service reply message with the client's public key. The hello request message is sent clear.

- 5 The hello reply messages are sent encrypted with the client's public key. All other messages are encrypted with the session key sent by the server in the HELLO ACK message. The HELLO ACK contains a session identifier allocated by the server (which must be sent with every succeeding message), and a session private key randomly generated by the server with which to encrypt all succeeding
- 10 messages.

```

//
enum SISNotifyCode
{
15   SIS_NTY_NOACCEPT = 101,    // server does not accept client
    connection
    SIS_NTY_SHUTDOWN = 102,    // server is shutting down
    SIS_NTY_TIMEOUT = 103,     // client has been idle too long
20   SIS_NTY_PROTOCOL = 104,   // invalid use of protocol
    SIS_NTY_SESSION = 105     // client session state failure
};

class SS_EXPORT_CLASS CSSSISNotifyMsg : public CSSSISMessage //
FINAL
25 {
public:
    SISNotifyCode      GetNotify () const { return m_notify; }
    const char*        GetDetail () const { return m_pDetail; }

30   virtual BOOL      Encode (SISNotifyCode,
                               const char* pDetail);

    virtual BOOL      Decode ();

35   virtual           CSSSISNotifyMsg ();
    virtual           ~CSSSISNotifyMsg ();

private:
40   SISNotifyCode      m_notify;
    char*              m_pDetail;
};

//
//
45 #define SIS_PROTO_VERSION 0x00010001

class SS_EXPORT_CLASS CSSSISHelloReqMsg : public CSSSISRequestMsg
50 {
public:
    unsigned long      GetProtoVersion () const { return m_protoVer; }
    unsigned char*     GetKeyData      ()      { return m_pKey; }
    DWORD             GetKeyLength    () const { return m_nKeyLen; }

```

```

    }

    virtual BOOL      Encode (SISSequence& sequence,
5                          unsigned char* pKey,
                          int nKeyLen);

    virtual BOOL      Decode ();

10    CSSSISHelloReqMsg ();
    ~CSSSISHelloReqMsg ();

private:
15    unsigned long      m_protoVer;
    unsigned char*      m_pKey;
    DWORD               m_nKeyLen;
};

20    class SS_EXPORT_CLASS CSSSISHelloACKMsg : public CSSSISACKMsg
    {
    public:
        SISSessionId      GetSessionId () const { return m_sessionId; }
        unsigned char*      GetKeyData () const { return m_pKey; }
        DWORD               GetKeyLength () const { return m_nKeyLen; }
25    virtual BOOL      Encode (SISSequence sequence,
        SISSessionId sessionId,
        unsigned char* pKey,
        int nKeyLen);
30    virtual BOOL      Decode ();

        CSSSISHelloACKMsg ();
        ~CSSSISHelloACKMsg ();
35    private:
        SISSessionId      m_sessionId;
        unsigned char*      m_pKey;
        DWORD               m_nKeyLen;
40    };
    //
45

```

The login message takes enough information to allow the user to log into the system. Generally, the login message is the second message sent (after HELLO). This message and all later messages are sent encrypted with the session key sent by the server in the HELLO ACK message.

50

The logoff message just specifies that the connection to the database is to be closed. This does not shut down the connection, but it does free all resources associated with that login. The commit parameter indicates whether the open transaction on the login should be committed or rolled back. If bCommit is

55 TRUE, then the transaction will be committed. Otherwise, it will be rolled back.

```

//
class SS_EXPORT_CLASS CSSSISLogonReqMsg : public CSSSISRequestMsg
{
5   public:
        const char*      GetInstance () const { return m_pInstance; }
        const char*      GetUserName () const { return m_pUserName; }
        const char*      GetPassword () const { return m_pPassword; }

10      virtual BOOL      Encode (SISSessionId sessionId,
                                SISSequence& sequence,
                                const char* pInstance,
                                const char* pUserName,
                                const char* pPassword);

15      virtual BOOL      Decode ();

                                CSSSISLogonReqMsg ();
                                ~CSSSISLogonReqMsg ();

20

private:
        char*            m_pInstance;
        char*            m_pUserName;
25      char*            m_pPassword;
};

class SS_EXPORT_CLASS CSSSISLogonACKMsg : public CSSSISACKMsg
{
30   public:
        virtual BOOL      Encode (SISSequence);

        virtual BOOL      Decode ();

35      CSSSISLogonACKMsg ();
      ~CSSSISLogonACKMsg ();

};

40  //
  //
class SS_EXPORT_CLASS CSSSISLogoffReqMsg : public CSSSISRequestMsg
{
45   public:
        BOOL              GetCommit () const { return m_bCommit; }

        virtual BOOL      Encode (SISSessionId sessionId,
                                SISSequence& sequence,
                                BOOL bCommit);

50      virtual BOOL      Decode ();

                                CSSSISLogoffReqMsg ();
                                ~CSSSISLogoffReqMsg ();

55

private:
        BOOL              m_bCommit;
60      };

class SS_EXPORT_CLASS CSSSISLogoffACKMsg : public CSSSISACKMsg
{
65   public:
        virtual BOOL      Encode (SISSequence);

```

```

        virtual BOOL        Decode ();
                                CSSSISLogoffACKMsg ();
                                ~CSSSISLogoffACKMsg ();
5    };

```

The SQL execution messages are intended to bundle communications into as few round-trips as possible. Typically, only one round trip is required for each user action. To accomplish this, several logical actions are performed by each request/reply. The major data structure exchanged over the protocol contains an array of values oriented into columns and rows (internally, the data is column major). Each column has a defined width and is guaranteed to be NUL terminated. All data are represented for transmission as strings. The data are encapsulated into the CSSSISColRowData which is used both to specify input bind variable values (in requests) and to return output bind variables (in ACKs).

To construct an instance of CSSSISColRowData, the structure is created by calling the function new. Then DefineCol is called for each column in the result set. The width must be specified at creation because the entire data block is allocated in one piece. Once all columns have been defined, Allocate is called to create the space and set up the pointers into the data block. The messages which take pointers to CSSSISColRowData blocks hang onto the pointers until deleted, but do not delete the passed-in blocks. Therefore, the blocks must be kept valid until the message which uses them is deleted, and the blocks must be separately deleted afterwards.

```

//
30    struct _SISColInfo;

    class SS_EXPORT_CLASS CSSSISColRowData
    {
    public:
35        BOOL        AddCol (int        nWidth,
                                const char* pName = NULL);
        BOOL        Allocate (int nRowCount);

        int        GetNumCols () const { return m_nCols; }
        int        GetNumRows () const { return m_nRows; }
40        int        GetColWidth (int        nColIndex) const;
        const char* GetColName (int        nColIndex) const;
        int        GetColIndex (const char* pName) const;
45        BOOL        GetColData (int        nColIndex,

```

```

char*      char* pBuf) const;
int         GetColDataPtr (int nColIndex);
            GetColDataSize (int nColIndex) const;

5  BOOL     GetColIndData (int nColIndex,
            short* pIndBuf) const;
short*     GetColIndDataPtr (int nColIndex);

10 char*     GetColRowData (int nColIndex,
            int nRowIndex);
short      GetColRowInd (int nColIndex,
            int nRowIndex) const;

15 void      Reset (void);

BOOL       SetColData (int nColIndex,
            const char* pBuf,
            int numRows);
20 BOOL     SetColIndData (int nColIndex,
            const short* pIndBuf,
            int numRows);

BOOL       SetColRowData (int nCol,
            int nRow,
            const char* pData);
25 BOOL     SetColRowInd (int nCol,
            int nRow,
            short indicator);

30 BOOL     IsSameShape (const CSSSISColRowData* pAs)
const;
BOOL       HasSameData (const CSSSISColRowData* pAs)
const;

35 void     Release ();

CSSSISColRowData ();
~CSSSISColRowData ();

40 private:
_SISColInfo* m_pCols;
int m_nCols;
int m_nColSize;
45 int m_nRows;
char* m_pDataBuf;
short* m_pIndBuf;
};

50 //

```

Before a statement can be executed (either DML or query), a "statement" must be opened. This may or may not correspond to an actual database cursor, but is used to specify the execution context for the statement (encapsulating whatever database state is necessary). For example, in an ODBC implementation, this corresponds to an HSTMT. The statement allocated this way must be closed with a FREESTMT request, otherwise the resources will be held open. The statement may be re-used for subsequent statements by sending an EXEC or a

QUERY message using the same statement. The EXEC request is used for a SQL statement which will not produce output columns (or where the client doesn't care) and requires only a single round trip. The given statement can be immediately reused for another statement as desired. Input bind variables can still be used with EXEC. In fact, a DML statement executed with EXEC can be executed multiple times by specifying multiple rows of data in the C\$\$\$ISColRowData block used for the input binds. All executions are done at the same time with a single round trip.

10 If the Commit flag is specified as TRUE, a commit operation will be done after the statement is executed. This helps reduce traffic by avoiding a separate COMMIT request after a single DML statement. If multiple DML statements are to be executed as a transaction, Commit should be specified only on the last one, or a separate COMMIT request should be made.

15

```

//
typedef unsigned long SISstmtId;
20  class SS_EXPORT_CLASS C$$$ISAllocStmtReqMsg : public C$$$ISRequestMsg
    {
    public:
        virtual BOOL      Encode (SISSessionId  sessionId,
25         SISSequence&    sequence);

        virtual BOOL      Decode ();

        C$$$ISAllocStmtReqMsg ();
        ~C$$$ISAllocStmtReqMsg ();
30    };

    class SS_EXPORT_CLASS C$$$ISAllocStmtACKMsg : public C$$$ISACKMsg
    {
    public:
35         SISstmtId      GetStmtId () const { return m_stmtId; }

        virtual BOOL      Encode (SISSequence  sequence,
        SISstmtId  stmt);
40         virtual BOOL      Decode ();

        C$$$ISAllocStmtACKMsg ();
        ~C$$$ISAllocStmtACKMsg ();

45    private:
        SISstmtId      m_stmtId;
    };

```

```

//
//
class SS_EXPORT_CLASS CSSSISExecReqMsg : public CSSSISRequestMsg
{
5 public:
    SISStmtId          GetStmtId () const { return m_stmtId; }
    const char*        GetStmt () const { return m_pStmt; }
    CSSSISColRowData*  GetInputBinds () { return m_pIBinds; }
    10    BOOL            GetCommit () const { return m_bCommit; }

    virtual BOOL        Encode (SISSessionId      sessionId,
                               SISSequence&       sequence,
                               SISStmtId          stmt,
                               const char*        pStmt,
                               15    CSSSISColRowData* pIBinds,
                               BOOL                bCommit);

    virtual BOOL        Decode ();

    20    CSSSISExecReqMsg ();
    ~CSSSISExecReqMsg ();

private:
    25    SISStmtId          m_stmtId;
    char*                  m_pStmt;
    CSSSISColRowData*      m_pIBinds;
    BOOL                   m_bCommit;
    30 };

class SS_EXPORT_CLASS CSSSISExecACKMsg : public CSSSISACKMsg
{
public:
    35    int                GetRowsAffected () const { return m_nRows; }

    virtual BOOL          Encode (SISSequence sequence,
                               int            nRowsAffected);

    40    virtual BOOL        Decode ();

    CSSSISExecACKMsg ();
    ~CSSSISExecACKMsg ();

    45 private:
        int                m_nRows;
    };

    50 //

```

The QUERY and FETCH requests are used when data is to be returned to the client (from a SELECT statement). Input binding, statement preparation, execution, and initial result fetching are all done by the query request. More data can be collected with the FETCH request. The QUERY request specifies the SQL statement, the input binding values as a CSSSISColRowData block, and the number of rows desired to be fetched immediately. The input binds may only specify a single row of values since the query will be executed only once per request. The reply will return no more than the number of rows specified. This

can be zero, in which case the query is executed, but no rows are fetched until a FETCH request is sent.

The QUERY request also defines the shape of the data to be retrieved. A
 5 second CSSSISColRowData block should be specified which has the columns
 defined, but no actual row values (Allocate never called). This is used by SISM
 in binding output values for retrieving results from the query. The output values
 in the QUERY and FETCH ACK messages will have the same shape as the block
 specified in the QUERY request. If the ACK returns all the rows available, the
 10 EOF member will be set to true. Further FETCH requests are unnecessary.

The FETCH request specifies a statement started by the QUERY request on
 which additional rows will be returned. As with QUERY, the reply will return no
 more than the number of rows specified as the maximum in the QUERY request.
 15 If the ACK returns all the rows available, the EOF member will be set to true.
 There is no need to make further FETCH requests.

```

//
20 class SS_EXPORT_CLASS CSSSISQueryReqMsg : public CSSSISRequestMsg
{
public:
    SISStmtId          GetStmtId () const { return m_stmtId; }
    const char*        GetStmt () const { return m_pStmt; }
    CSSSISColRowData*  GetInputBinds () { return m_pIBinds; }
25 CSSSISColRowData*  GetOutputDesc () { return m_pODesc; }
    int                GetMaxFetch () const { return m_nMaxFetch; }

    virtual BOOL       Encode (SISSessionId      sessionId,
                               SISSequence&      sequence,
30                               SISStmtId        stmt,
                               const char*       pStmt,
                               CSSSISColRowData* pIBinds,
                               CSSSISColRowData* pODesc,
                               int                nMaxFetch);

35    virtual BOOL       Decode ();

                               CSSSISQueryReqMsg ();
                               ~CSSSISQueryReqMsg ();

40
private:
    SISStmtId          m_stmtId;
    char*              m_pStmt;
    CSSSISColRowData*  m_pIBinds;
45 CSSSISColRowData*  m_pODesc;
    int                m_nMaxFetch;
};

```

```

class SS_EXPORT_CLASS CSSSISQueryACKMsg : public CSSSISACKMsg
{
public:
    CSSSISColRowData* GetFetchVals () { return m_pFetch; }
    int GetActualRows () const { return m_nActual; }
    5     BOOL GetEOF () const { return m_bEOF; }

    virtual BOOL Encode (SISSequence sequence,
                        CSSSISColRowData* pFetch,
                        10     int nActual,
                        BOOL bEOF);

    virtual BOOL Decode ();

    15     CSSSISQueryACKMsg ();
    ~CSSSISQueryACKMsg ();

private:
    20     CSSSISColRowData* m_pFetch;
    int m_nActual;
    BOOL m_bEOF;
};

//
//
class SS_EXPORT_CLASS CSSSISFetchReqMsg : public CSSSISRequestMsg
{
public:
    30     SISStmtId GetStmtId () const { return m_stmtId; }

    virtual BOOL Encode (SISSessionId sessionId,
                        SISSequence& sequence,
                        SISStmtId stmt);

    35     virtual BOOL Decode ();

    CSSSISFetchReqMsg ();
    ~CSSSISFetchReqMsg ();

    40

private:
    SISStmtId m_stmtId;
    45 };

class SS_EXPORT_CLASS CSSSISFetchACKMsg : public CSSSISACKMsg
{
public:
    CSSSISColRowData* GetFetchVals () { return m_pFetch; }
    50     int GetActualRows () const { return m_nActual; }
    BOOL GetEOF () const { return m_bEOF; }

    virtual BOOL Encode (SISSequence sequence,
                        CSSSISColRowData* pFetch,
                        55     int nActual,
                        BOOL bEOF);

    virtual BOOL Decode ();

    60     CSSSISFetchACKMsg ();
    ~CSSSISFetchACKMsg ();

private:
    65     CSSSISColRowData* m_pFetch;
    int m_nActual;

```

```

        BOOL                m_bEOF;
    };
    //
5

```

10 The FREESTMT request closes a statement created with the ALLOCSTMT request. This implicitly invalidates any query which is currently active on this statement. The COMMIT request commits outstanding DML operations on the current session (login). The ROLLBACK request rolls back outstanding DML operations on the current session (login).

```

15 //
class SS_EXPORT_CLASS CSSSISFreeStmtReqMsg : public CSSSISRequestMsg
{
public:
    SISStmtId                GetStmtId () const { return m_stmtId; }
20     virtual BOOL           Encode (SISSessionId sessionId,
                                   SISSequence& sequence,
                                   SISStmtId stmt);

25     virtual BOOL           Decode ();

                                   CSSSISFreeStmtReqMsg ();
                                   ~CSSSISFreeStmtReqMsg ();

30 private:
    BOOL                    m_stmtId;
};

35 class SS_EXPORT_CLASS CSSSISFreeStmtACKMsg : public CSSSISACKMsg
{
public:
    virtual BOOL             Encode (SISSequence);
40     virtual BOOL           Decode ();

                                   CSSSISFreeStmtACKMsg ();
                                   ~CSSSISFreeStmtACKMsg ();

};

45 //
//
class SS_EXPORT_CLASS CSSSISCommitReqMsg : public CSSSISRequestMsg
50 {
public:
    virtual BOOL             Encode (SISSessionId sessionId,
                                   SISSequence& sequence);

55     virtual BOOL           Decode ();

                                   CSSSISCommitReqMsg ();
                                   ~CSSSISCommitReqMsg ();

};

```

```

class SS_EXPORT_CLASS CSSSISCommitACKMsg : public CSSSISACKMsg
{
public:
5   virtual BOOL      Encode (SISSequence);
   virtual BOOL      Decode ();
   CSSSISCommitACKMsg ();
   ~CSSSISCommitACKMsg ();
10  };
   //
   //
15  class SS_EXPORT_CLASS CSSSISRollbackReqMsg : public CSSSISRequestMsg
   {
public:
   virtual BOOL      Encode (SISSessionId sessionId,
20                            SISSequence& sequence);
   virtual BOOL      Decode ();
   CSSSISRollbackReqMsg ();
   ~CSSSISRollbackReqMsg ();
25  };
class SS_EXPORT_CLASS CSSSISRollbackACKMsg : public CSSSISACKMsg
{
public:
30   virtual BOOL      Encode (SISSequence);
   virtual BOOL      Decode ();
   CSSSISRollbackACKMsg ();
   ~CSSSISRollbackACKMsg ();
35  };
   //
40

```

This next group of messages support remote file access through the
 SISMGR, which is used to implement literature and correspondence (fulfillment)
 requests. These messages provide simple file-style access to files located on the
 machine running SISMGR or remotely accessible from it. The basic file
 45 operations are supported: open, read, write and close.

Files are requested by category and name. The category can be used to
 identify different groups of documents and to allow the files to be partitioned for
 different sets of users. The server may also enforce different access restrictions
 50 on different categories. An open file is identified by a SISFileId which is returned
 by open and which all other file-related requests must provide. This file ID must
 be closed when the operations are complete.

```

//
typedef unsigned long SISFileId;

enum SISFileMode
5 {
    SIS_MODE_NONE      = 0,
    SIS_MODE_READ       = (1 << 0),
    SIS_MODE_WRITE      = (1 << 1),
    SIS_MODE_RDWR       = SIS_MODE_READ | SIS_MODE_WRITE
10 };

class SS_EXPORT_CLASS CSSSISFileInfoReqMsg : public CSSSISRequestMsg
{
public:
15     const char*      GetCategory () const { return m_pCategory; }
    const char*      GetFileName () const { return m_pFileName; }

    virtual BOOL      Encode (SISSessionId  sessionId,
20                          SISSequence&  sequence,
                          const char*    pCategory,
                          const char*    pFileName);

    virtual BOOL      Decode ();

25     CSSSISFileInfoReqMsg ();
    ~CSSSISFileInfoReqMsg ();

private:
30     char*            m_pCategory;
    char*            m_pFileName;
};

#ifdef _TIME_T_DEFINED
35     typedef long     time_t;
#endif

typedef unsigned int  size_t;

40     class SS_EXPORT_CLASS CSSSISFileInfoACKMsg : public CSSSISACKMsg
    {
    public:
        const char*      GetFileName () const { return m_pFileName; }
45        size_t          GetFileSize () const { return m_nFileSize; }
        time_t           GetCreateTime () const { return m_createTime; }
        time_t           GetWriteTime () const { return m_writeTime; }
50        SISFileMode     GetFileModes () const { return m_modes; }

        virtual BOOL      Encode (SISSequence sequence,
55                          const char* pFileName,
                          size_t      nFileSize,
                          time_t      createTime,
                          time_t      writeTime,
                          SISFileMode modes);

60        virtual BOOL      Decode ();

        CSSSISFileInfoACKMsg ();
        ~CSSSISFileInfoACKMsg ();

65    private:

```

```

        char*                m_pFileName;
        unsigned long        m_nFileSize;
        time_t               m_createTime;
        time_t               m_writeTime;
5      SISFileMode          m_modes;
    };

10  class SS_EXPORT_CLASS CSSSISFileOpenReqMsg : public CSSSISRequestMsg
    {
    public:
        const char*          GetCategory () const { return m_pCategory; }
        const char*          GetFileName () const { return m_pFileName; }
15      SISFileMode          GetOpenMode () const { return m_openMode; }

        virtual BOOL         Encode (SISSessionId sessionId,
                                   SISSequence& sequence,
                                   const char* pCategory,
20      const char* pFileName,
                                   SISFileMode openMode);

        virtual BOOL         Decode ();

25      CSSSISFileOpenReqMsg ();
        ~CSSSISFileOpenReqMsg ();

    private:
        char*                m_pCategory;
        char*                m_pFileName;
30      SISFileMode          m_openMode;
    };

35  class SS_EXPORT_CLASS CSSSISFileOpenACKMsg : public CSSSISACKMsg
    {
    public:
        SISFileId            GetFileId () const { return m_fileId; }
        const char*          GetFileName () const { return m_pFileName; }
40      size_t               GetFileSize () const { return m_nFileSize; }
        time_t               GetCreateTime () const { return m_createTime; }
        time_t               GetWriteTime () const { return m_writeTime; }
45      }

        virtual BOOL         Encode (SISSequence sequence,
                                   SISFileId file,
                                   const char* pFileName,
50      size_t nFileSize,
                                   time_t createTime,
                                   time_t writeTime);

        virtual BOOL         Decode ();

55      CSSSISFileOpenACKMsg ();
        ~CSSSISFileOpenACKMsg ();

60  private:
        SISFileId            m_fileId;
        char*                m_pFileName;
        unsigned long        m_nFileSize;
65      time_t               m_createTime;
        time_t               m_writeTime;

```



```

};

5  class SS_EXPORT_CLASS CSSSISFileReadReqMsg : public CSSSISRequestMsg
    {
    public:
        SISFileId      GetFileId      () const { return m_fileId; }
        size_t          GetChunkSize   () const { return m_nChunkSize; }
10     int              GetMaxChunks   () const { return m_nMaxChunks; }
    }

        virtual BOOL    Encode (SISSessionId sessionId,
15                               SISSequence& sequence,
                               SISFileId      file,
                               size_t         chunkSize,
                               int            maxChunks);

        virtual BOOL    Decode ();

20     CSSSISFileReadReqMsg ();
        ~CSSSISFileReadReqMsg ();

25     private:
        SISFileId      m_fileId;
        size_t          m_nChunkSize;
        int             m_nMaxChunks;
    };

30     class SS_EXPORT_CLASS CSSSISFileReadACKMsg : public CSSSISACKMsg
        {
        public:
            unsigned char* GetChunkData () const { return m_pData; }
35             size_t      GetChunkSize () const { return m_nDataLen; }
            BOOL          GetEOF       () const { return m_bEOF; }

            virtual BOOL    Encode (SISSequence sequence,
40                                   unsigned char* pData,
                                   size_t      nDataLen,
                                   BOOL        bEOF);

            virtual BOOL    Decode ();

45     CSSSISFileReadACKMsg ();
        ~CSSSISFileReadACKMsg ();

        private:
            unsigned char* m_pData;
50             size_t      m_nDataLen;
            BOOL          m_bEOF;
        };

55     class SS_EXPORT_CLASS CSSSISFileWriteReqMsg : public CSSSISRequestMsg
        {
        public:
            SISFileId      GetFileId      () const { return m_fileId; }
            unsigned char* GetChunkData   () const { return m_pData; }
60             size_t      GetChunkSize   () const { return m_nDataLen; }
        }

        virtual BOOL    Encode (SISSessionId sessionId,
65                               SISSequence& sequence,
                               SISFileId      file,
                               unsigned char* pData,

```

```

                                size_t      nDataLen);

        virtual BOOL      Decode ();

5         CSSSISFileWriteReqMsg ();
        ~CSSSISFileWriteReqMsg ();

10      private:
        SISFileId          m_fileId;
        unsigned char*     m_pData;
        size_t             m_nDataLen;
    };

15      class SS_EXPORT_CLASS CSSSISFileWriteACKMsg : public CSSSISACKMsg
    {
    public:
        virtual BOOL      Encode (SISSequence sequence);

20         virtual BOOL      Decode ();

        CSSSISFileWriteACKMsg ();
        ~CSSSISFileWriteACKMsg ();

25     };

    class SS_EXPORT_CLASS CSSSISFileCloseReqMsg : public CSSSISRequestMsg
    {
    public:
30         SISFileId          GetFileId () const { return m_fileId; }

        virtual BOOL      Encode (SISSessionId sessionId,
                                SISSequence& sequence,
                                SISFileId file);

35         virtual BOOL      Decode ();

        CSSSISFileCloseReqMsg ();
        ~CSSSISFileCloseReqMsg ();

40     private:
        SISFileId          m_fileId;
    };

45      class SS_EXPORT_CLASS CSSSISFileCloseACKMsg : public CSSSISACKMsg
    {
    public:
50         virtual BOOL      Encode (SISSequence sequence);

        virtual BOOL      Decode ();

        CSSSISFileCloseACKMsg ();
        ~CSSSISFileCloseACKMsg ();

55     };
    //

```

Following are wrapper functions for creating all of the different classes
 60 implemented in this library, if SISNAPI is not implemented as a DLL.

```

//
extern BOOL SS_EXPORT_API SISNewMessage (SISMessageType msgType,

```

```

SISRequestCode reqCode,
CSSSISMessage*& pReturn);

5  extern BOOL SS_EXPORT_API SISNewMessage (SISMessageType msgType,
                                           CSSSISMessage*& pReturn);

extern BOOL SS_EXPORT_API SISNewColRowData (CSSSISColRowData*&
pReturn);

10  extern BOOL SS_EXPORT_API SISNewErrorInfo (CSSSISErrorInfo*&
pReturn);

15  #endif // __SISNAPI_H__

Configuration
SISMGR creates database connections in response to client requests
20  and actions, and does not create an initial connection to the
corporate database. At startup, SISMGR loads its configuration
information from the Windows NT registry. Configuration information
is stored as values set on the Configuration key. The configuration
parameters read at startup are as follows:

25  • Allow Connections From
Allows specification of a set of host names or network addresses from
which connections will be accepted. Network addresses are specified
in the usual format and allow sub-nets to be specified with
30  wildcarding (for example: 206.79.143.* for all addresses in the
206.79.143 "class C" subnet). Multiple host names and network addresses are
separated by commas. If specified as *.*.*.*, no host address validation is
performed.

• Internet Port Number
35  Specifies the Internet port number at which the server will listen for client
connections. Clients must connect to the correct machine and to this port number
to contact SISMGR.

• Maximum Connections
Specifies the maximum number of connections allowed at any one time. If
40  specified as 0, no limit is placed on the number of connections.

• Maximum Idle Minutes
Specifies the maximum number of minutes a client can be idle (not sending any
messages) before the connection is automatically closed by SISMGR. This can be
used to free up resources held by a client when a salesperson leaves it running. If

```

specified as 0, clients are never timed out.

- **Connect String Map**

Specifies the mappings between "connect strings" specified by the client when connecting and database vendors and "instance" identifiers for instantiating a connector when the client logs in. The format of this value is a sequence of connect string "maps" separated by commas. Each map is comprised of the connect string as specified by the user, the database vendor this corresponds to, and the database instance to connect to. For example: MPengr=oracle:dev indicates that the user-supplied connect string "MPengr" will be implemented by the Oracle connector using the database instance "dev" (which must be known to the machine on which SISMGR is running).

- **Oracle Connector DLL**

Specifies the name of the DLL which implements the connector to Oracle databases. This will be used for connections which map to Oracle database instances. The value should be the name of the DLL which implements the SISM connector for Oracle. For example: ssx30o72.dll for the Oracle 7.2 connector.

- **Sybase Connector DLL**

Specifies the name of the DLL which implements the connector to Sybase SQL Server databases. This will be used for connections which map to Oracle database instances. The value should be the name of the DLL which implements the SISM connector for Sybase SQL Server. For example: ssx30s11.dll for the Sybase system 11 connector.

- **Informix Connector DLL**

Specifies the name of the DLL which implements the connector to Informix On-line databases. This will be used for connections which map to Informix database instances. The value should be the name of the DLL which implements the SISM connector for Informix On-line. For example: ssx30i72.dll for the Informix 7.2 connector.

- **Microsoft Connector DLL**

Specifies the name of the DLL which implements the connector to Microsoft SQL Server databases. This will be used for connections which map to Microsoft

database instances. The value should be the name of the DLL which implements the SISM connector for Microsoft SQL Server. For example: `ssx30m60.dll` for the Microsoft 6.0 connector.

- **Update Frequency**

- 5 Specifies the frequency in seconds with which SISMGR updates the statistics in the registry key for monitoring programs and checks the control key for shutdown requests. The default is 5 seconds, which allows for fine-grained monitoring.

- **Message Resource DLL**

- 10 Specifies the name of the DLL which is loaded for handling error, warning, and informational messages. Internally, error codes are used and are translated into strings only when errors are reported to the user. For example, `ssrsis30.enu` for US English.

- **File Root Directory**

- 15 Specifies the absolute path name to the directory which contains the files which are available to be served to clients using the file access messages.

- **Configuration Password**

- 20 Specifies the password which must be entered by the user when starting up the configuration user interface applet. This value is not read or set by SISMGR itself.

A user interface may be provided for editing the registry values. For example, a Windows NT control panel applet may be used, or some other form to make the UI consistent with other system services.

- 25 Once the system preferences have been read, SISMGR is up and begins accepting connections.

CONCLUSION

- 30 Various modifications to these embodiments will be readily apparent to those skilled in the art, and the generic principles defined herein may be applied to other embodiments without the use of inventive faculty. Thus, the present invention is not intended to be limited to the embodiments shown herein, but is to

be accorded the widest scope consistent with the principles and novel features disclosed herein.

5 All publications and patent applications mentioned in this specification are herein incorporated by reference to the same extent as if each individual publication or patent application was specifically and individually indicated to be incorporated by reference.

The invention now being fully described, it will be apparent to one of ordinary skill in the art that many changes and modifications can be made thereto without departing therefrom.

10

APPENDIX

Writing User Transaction Log File for a Given Laptop Node

This program will be called by a server-side process that processes transaction log entries for all Laptop Nodes. For each Laptop Node, the calling process is building the UserTrxnLogFileName and calling Program 1.

Input Parameters

- LaptopNodeId - node_id of the destination laptop
- UserTrxnLogFileName - full path of the file where txns will be written
- MaxBatchTxns - number of txns between commits and updates to the S_DOCK_STATUS table
- MaxTxns - number of txns to process in this session. Use this parameter to limit processing.

15 MAIN ALGORITHM

```

-- CHECK PARAMETERS
IF (MAXTXNS < 1 || MAXBATCHTXNS < 1) THEN
  INVALID PARAMETER
20  END IF

-- GET LAST LOG_EXTRACT NUMBER FOR THE LAPTOP FROM S_DOCK_STATUS
LAST_TXN_COMMIT_NUMBER = UTLDSTATGETLOGNUM(LAPTOPNODEID);

25  -- INITIALIZE VARIABLES
  NUMTXNS = 0;          -- TOTAL NUMBER OF TXNS PROCESSED
  NUMBATCHTXNS = 0;    -- TOTAL NUMBER OF TXNS WRITTEN IN THE CURRENT
  BATCH

30  -- READ DOCKING OBJECT AND TABLE DEFINITIONS INTO MEMORY STRUCTURES
  STARTDICTAPI ();

  -- OPEN THE USER LOG TXN FILE
  OPEN USER LOG TXN FILE

35  -- SELECT AND PROCESS NEW TXNS IN S_DOCK_TRANSACTION_LOG
  -- WHERE TXN_COMMIT_NUMBER > LAST_TXN_COMMIT_NUMBER
  FOR EACH NEW TXN LOOP

40    -- STOP PROCESSING IF REACH MAXTXNS
    IF NUMTXNS = MAXTXNS THEN
      BREAK;
    END IF;

45    -- PREVENT CIRCULAR TXNS. DO NOT SEND THE TXN BACK TO THE
    ORIGINATING LAPTOP
    IF TXN.ORIGINNODEID = LAPTOPNODEID THEN
      GOTO NEXT TRANSACTION
    END IF;

50    -- PROCESS ALL OTHER TYPES OF TRANSACTIONS

    -- THIS IS THE VISIBILITY CALCULATOR!
    -- THIS ROUTINE ALSO PROCESSES IMPLICIT VISIBILITY EVENTS
55    -- LATER: DATA MERGE CAN CALL THIS FUNCTION TO CHECK WHETHER A TXN
    IS
    -- STILL VISIBLE WHEN MERGING TXNS INTO A LAPTOP OR SERVER
    DATABASE.

```

```

    CHECKVISIBILITY (LAPTOPNODEID, LOGRECORDTYPE, TABLENAME,
TRANSROWID);
    IF TXN IS VISIBLE THEN
      -- WRITE TRANSACTIONS TO USERTXNLOG FILE DEPENDING ON THE
      -- TYPE OF LOGRECORDTYPE.
      WRITE THE TXN TO THE USER LOG FILE
      ++NUMBATCHTXNS
    END IF;

10  ---FINISHED PROCESSING THE TXN
    -- COMMIT (IF NEEDED)
    IF NUMBATCHTXNS = MAXBATCHTXNS THEN
      -- ASSUME THAT SEPARATE PROCESS COMES AROUND AND DELETES
      -- TXNS IN S_DOCK_TRANSACTION_LOG THAT HAVE BEEN PROCESSED
15  -- FOR ALL NODES. SO, NO NEED TO DELETE THE TXNS FROM THE LOG.
      UPDATE LAST LOG_EXTRACT NUMBER FOR LAPTOP IN S_DOCK_STATUS
      COMMIT;
      NUMBATCHTXNS = 0
    END IF;

20  ++NUMTXNS
    END LOOP; /* EACH TRANSACTION IN THE TXN LOG TABLE */

    -- COMMIT
25  UPDATE LAST LOG_EXTRACT NUMBER FOR LAPTOP IN S_DOCK_STATUS
    COMMIT;

    -- CLOSE LOG FILE (IF NEEDED)
    IF USERTXNLOGFILEP != NULL THEN
30  CLOSE FILE;
    END IF;

    STOPDICTAPI ();

35  CHECK VISIBILITY ROUTINES

    -- CHECK IF A RECORD IN THE TXN LOG IS VISIBLE TO A LAPTOPNODEID
    BOOL CHECKVISIBILITY (LAPTOPNODEID, LOGRECORDTYPE, TABLENAME,
TRANSROWID)
40  {
    -- SQLSTATEMENTS ROUTED BASED ON THE DESTINATION LIST
    IF LOGRECORDTYPE IN ('SQLSTATEMENT') THEN
      IF LAPTOP NODE IN DESTINATION LIST THEN
        RETURN TRUE;
45  END IF;

    -- SHADOW AND MULTI RECORD LOGRECORDTYPES ARE ROUTED TO ALL NODES
    -- NO VISIBILITY EVENTS WITH THESE LOGRECORDTYPES.
    ELSIF LOGRECORDTYPE IN ('SHADOWOPERATION', 'MULTIRECORDDELETE',
50  'MULTIRECORDUPDATE') THEN
      RETURN TRUE;

    -- SIMPLE DELETES NEED MORE PROCESSING
    ELSIF LOGRECORDTYPE IN ('SIMPLE DELETE') THEN
55  IF (TABLE.VISIBILITY IN ('ENTERPRISE', 'LIMITED')) THEN
      RETURN TRUE;
    END IF;

    -- SIMPLE INSERTS AND SIMPLE UPDATES NEED MORE PROCESSING
    -- CHECKTXNVISIBILITY() ALSO PROCESSES IMPLICIT VISIBILITY EVENTS
60  ELSIF LOGRECORDTYPE IN ('SIMPLE INSERT', 'SIMPLE UPDATE') THEN
      IF (TABLE.VISIBILITY = 'ENTERPRISE') THEN
        RETURN TRUE;
      ELSIF TABLE.VISIBILITY = 'LIMITED' THEN
65  IF CHECKTXNVISIBILITY (LAPTOPNODEID, TABLE, ROWID) THEN

```



```

        RETURN TRUE;
      END IF;
    END IF;
  END IF;
5  }

  -- CHECK IF A RECORD IN THE TXN LOG IS VISIBLE TO A LAPTOPNODEID
  STATIC BOOL CHECKTXNVISIBILITY (LAPTOPNODEID, TABLE, ROWID)
  {
10    BOOL BVISIBLE = FALSE;

    FIND THE TABLE IN THE DICTIONARY;
    IF TABLE NOT FOUND THEN
      ERROR: TABLE NOT DEFINED
15    END IF;

    FOR ALL DOCKING OBJECTS THAT THE TABLE BELONGS TO LOOP
      -- GENERATE SQL TO GET PRIMARYID VALUES OF THE DOCKING OBJECT
      GENERATEPRIMARYIDSQ (TABLE, ROWID, DOCKINGOBJECT);
20    FOR EACH PRIMARYID VALUE RETRIEVED LOOP
      CHECKOBJECTVISIBILITY (LAPTOPNODEID, PRIMARYTABLE,
        PRIMARYROWID)
      IF OBJECT IS VISIBLE THEN
        -- BECAUSE CHECKOBJECTVISIBILITY() ALSO PROCESSES IMPLICIT
25        -- VISIBILITY EVENTS, WE MUST LOOP THROUGH ALL DOCKING OBJECTS
        -- EVEN IF WE ALREADY KNOW THAT THE TXN IS VISIBLE.
        -- EXCEPTION: IF THE TABLE HAS VIS_EVENT_FLG = 'N'
        -- THEN WE CAN RETURN IMMEDIATELY.
        IF TABLE.VISIBILITYEVENTFLG = 'N' THEN
30          RETURN TRUE;
        ELSE
          BVISIBLE = TRUE;
        END IF;
      END IF;
35    END LOOP;
    END LOOP;

    RETURN BVISIBLE;
40  }

  -- CHECK IF AN INSTANCE OF A DOCKING OBJECT IS VISIBLE TO THE LAPTOP
  USER.
  -- ALSO PROCESSES IMPLICIT VISIBILITY EVENTS!
45  BOOL CHECKOBJECTVISIBILITY (LAPTOPNODEID, DOCKINGOBJECTNAME,
    PRIMARYROWID)
  {
    FOR EACH VISIBILITY RULE FOR THE DOCKING OBJECT LOOP
      IF RULETYPE = RULESQL THEN
50        RUN THE SELECT SQL STATEMENT USING PRIMARYROWID;
        IF ANY ROWS RETURNED THEN
          -- ROW IS VISIBLE
          -- PROCESS AN IMPLICIT DOWNLOAD OBJECT
          DOWNLOADOBJECTINSTANCE (LAPTOPNODEID, PRIMARYTABLENAME,
55            PRIMARYROWID);
          RETURN TRUE;
        END IF;
      ELSIF RULETYPE = CHECKDOCKINGOBJECT THEN
        RUN THE PARAMETERSQ (LAPTOPNODEID, PRIMARYROWID) TO GET NEWPRIMARYROWID
60        FOR EACH RECORD RETRIEVED BY PARAMETERSQ LOOP
          -- RECURSIVE!
          CHECKOBJECTVISIBILITY (LAPTOPNODEID, CHECKDOCKINGOBJECTNAME,
            NEWPRIMARYROWID);

          IF RC = TRUE THEN
65            -- PROCESS AN IMPLICIT DOWNLOAD OBJECT
            DOWNLOADOBJECTINSTANCE (LAPTOPNODEID, PRIMARYTABLENAME,

```

```

                                PRIMARYROWID);
                                RETURN TRUE;
                                END IF;
                                END LOOP;
5      END IF;
      END LOOP;

      -- OBJECT IS NOT VISIBLE.

10     -- PROCESS AN IMPLICIT REMOVE OBJECT
      REMOVEOBJECTINSTANCE (LAPTOPNODEID, PRIMARYTABLENAME,
      PRIMARYROWID);

15     RETURN FALSE;
  }

  GENERATE SQL STATEMENT TO GET PRIMARYID

  -- GENERATE THE SELECT SQL STATEMENT TO GET THE PRIMARYID VALUE OF
20  -- THE DOCKING OBJECT FOR THE GIVEN MEMBERTABLE
  --
  -- SQL STATEMENT LOOKS LIKE:
  --   SELECT TP.<ROW_ID>
  --     FROM <TABLE_OWNER>.<TABLE> T1,
25  --        <TABLE_OWNER>.<PKTABLE> T2,
  --        ... ONE OR MORE INTERMEDIATE TABLES BETWEEN THE TABLE
  --        AND THE PRIMARYTABLE
  --        <TABLE_OWNER>.<PKTABLE> TN
  --        <TABLE_OWNER>.<PRIMARYTABLE> TP
30  --   WHERE T1.ROW_ID = :ROW_ID /* ROW_ID IN TRANSACTION LOG */
  --        /* JOIN TO PK TABLE T2 */
  --        AND T1.<FKCOLUMN> = T2.<PKCOLUMN>
  --        AND <T1 FKCONDITION>
  --        /* ANY NUMBER OF JOINS UNTIL REACH THE TABLE THAT JOINS
35  --        TO THE PRIMARYTABLE */
  --        /* JOIN FROM T2 TO TN */
  --        AND T2.<FKCOLUMN> = TN.<PKCOLUMN>
  --        AND <T2 FKCONDITION>
  --        /* JOIN TO THE PRIMARYTABLE */
40  --        AND TN.<FKCOLUMN> = TP.<PKCOLUMN>
  --        AND <TN FKCONDITION>
  --
  -- NOTE THAT THERE MAY BE ONE OR MORE PATHS FROM THE MEMBER TABLE
  -- TO THE PRIMARY TABLE. WE NEED TO GENERATE A SQL SELECT STATEMENT
45  -- FOR EACH OF THE PATHS AND UNION THE STATEMENTS TOGETHER.
  --
  -- THIS FUNCTION ASSUMES THAT THERE ARE NO LOOPS IN THE DEFINITION.
  --
  -- THESE SQL STATEMENT DO NOT CHANGE FOR EACH TABLE IN A DOCKING
50  -- OBJECT,
  -- SO WE CAN CALCULATE THEM ONE AND STORE THEM IN MEMORY.
  --
  STRUCT
  {
55    CHAR* SELECTLIST;
    CHAR* FROMCLAUSE;
    CHAR* WHERECLAUSE;
    UINT NUMTABLES; /* ALSO THE NUMBER OF JOINT TO REACH THE PRIMARY
60    TABLE */
  } GENSTMT;

  GENERATEPRIMARYIDSQ (TABLE, DOCKINGOBJECT)
65  {
    /* THERE MAY BE MORE THAN ONE SQL STATEMENT, SO WE HAVE A DYNAMIC

```

```

        ARRAY OF SQL STATEMENTS. EACH ELEMENT IN THE ARRAY IS A PATH
        FROM THE TABLE TO THE PRIMARY TABLE*/
        DYNARRID GENSTMTARR;
        GENSTMT NEWGENSTMT;

5      CHAR* SQLSTMT;

        DYNARRCREATE (GENSTMTARR);

10     -- CREATE THE FIRST ELEMENT AND INITIALIZE
        NEWGENSTMT = MALLOC();
        NEWGENSTMT.NUMTABLES = 1;
        NEWGENSTMT.SELECTLIST = "SELECT ROW_ID";
        NEWGENSTMT.FROMCLAUSE = "FROM <TABLE> T1";
15     NEWGENSTMT.WHERECLAUSE = "WHERE T1.ROW_ID = :ROW_ID";
        DYNARRAPPEND (GENSTMTARR, &NEWGENSTMT);

        /* RECURSIVELY FOLLOW FKS TO THE PRIMARYTABLE */
        BUILD THE SELECT, FROM AND WHERE CLAUSE SIMULTANEOUSLY */
20     ADDPKTABLE (TABLE, DOCKINGOBJECT, GENSTMTARR, 0);

        -- UNION ALL THE PATHS TOGETHER
        NUMSTMTS = DYNARRSIZE (GENSTMTARR);
        FOR ALL ELEMENTS IN THE ARRAY LOOP
25     TMSQLSTMT = GENSTMTARR[J].SELECTLIST || GENSTMTARR[J].FROMCLAUSE
        || GENSTMTARR[J].WHERECLAUSE;
        SQLSTMT = SQLSTMT || 'UNION' || TMSQLSTMT;
        END LOOP;

30     DYNARRDESTROY (GENSTMTARR);

        IF SQLSTMT = NULL THEN
            ERROR: NO PATH FROM TABLE TO PRIMARY TABLE.
        END IF;
35 }

        -- RECURSIVELY FOLLOW ALL FKS TO THE PRIMARY TABLE
        ADDPKTABLE (TABLE, DOCKINGOBJECT, GENSTMT, INPUTSTMTNUM)
40     {
        UINT NUMFKS = 0;
        UINT STMTNUM;
        GENSTMT NEWGENSTMT;

45     FOR ALL FKS FOR THE TABLE LOOP
        IF PKTABLE IS A MEMBER TABLE OF THE DOCKING OBJECT THEN
            -- IF THERE'S MORE THAN ONE FK, THEN THERE IS MORE THAN ONE PATH
            -- OUT OF THE CURRENT TABLE.
            -- COPY THE SQL STMT TO A NEW DYNARRELMT TO CREATE A NEW PATH
50     IF NUMFKS > 0 THEN
            -- CREATE A NEW ELEMENT AND COPY FROM GENSTMT[INPUTSTMTNUM]
            NEWGENSTMT = MALLOC();
            NEWGENSTMT.NUMTABLES = GENSTMT[INPUTSTMTNUM].NUMTABLES;
            NEWGENSTMT.SELECTLIST = GENSTMT[INPUTSTMTNUM].SELECTLIST;
55     NEWGENSTMT.FROMCLAUSE = GENSTMT[INPUTSTMTNUM].FROMCLAUSE;
            NEWGENSTMT.WHERECLAUSE = GENSTMT[INPUTSTMTNUM].WHERECLAUSE;
            DYNARRAPPEND (GENSTMTARR, &NEWGENSTMT);
            STMTNUM = DYNARRSIZE (GENSTMTARR);

60     -- PUT A CHECK HERE FOR INFINITE LOOPS
            IF STMTNUM == 20 THEN
                ERROR: PROBABLY GOT AN INFINITE LOOP?
            END IF;
            ELSE
65     STMTNUM = INPUTSTMTNUM;
            END IF;

```

```

-- APPEND THE NEW PKTABLE TO THE FROMCLAUSE AND WHERECLAUSE
GENSTMT[STMTNUM].FROMCLAUSE =
    GENSTMT[STMTNUM].FROMCLAUSE || ",\N <TABLE> T<NUMTABLES +
5 1>";
    GENSTMT[STMTNUM].WHERECLAUSE =
    GENSTMT[STMTNUM].WHERECLAUSE ||
    "AND T<NUMTABLES>.<FKCOLUMN> = T<NUMTABLES + 1>.<PKCOLUMN>"
    ||
10 "AND <FKCONDITION FOR TABLE IF ANY>";
    ++GENSTMT.NUMTABLES;

    -- PKTABLE IS THE PRIMARY TABLE THEN DONE.
    IF PKTABLE = PRIMARYTABLE THEN
        RETURN;
15 ELSE
    ADDPKTABLE (PKTABLE, DOCKINGOBJECT, GENSTMT, STMTNUM);
    END IF;

    -- ONLY COUNT FKS TO OTHER MEMBER TABLES IN THE SAME DOCKING
20 OBJECT
    ++NUMFKS;

    END IF;
    END LOOP;
25 RETURN;
}

PROCESS VISIBILITY EVENTS

30 -- DOWNLOAD AN OBJECT INSTANCE TO A LAPTOP
-- THIS FUNCTION ALSO DOWNLOADS ALL RELATED DOCKING OBJECT INSTANCES.
BOOL DOWNLOADOBJECTINSTANCE (LAPTOPNODEID, OBJECTNAME, PRIMARYROWID)
{
35 -- CHECK IF THE OBJECT INSTANCE IS ALREADY DOWNLOADED TO THE LAPTOP
    FIND THE OBJECT INSTANCE IN THE S_DOBJ_INST TABLE
    IF EXISTS ON LAPTOP THEN
        RETURN TRUE;
    END IF;

40 -- REGISTER OBJECT INSTANCE IN S_DOBJ_INST TABLE

    -- WRITE DOWNLOAD OBJECT RECORDS TO THE TXN LOG
    FOR EACH MEMBER TABLE OF THE DOCKING OBJECT LOOP
        GENERATE SQL SELECT STATEMENT TO DOWNLOAD RECORDS
45 WRITE EACH RETRIEVED RECORD TO THE USER TXN LOG FILE
    END LOOP;

    -- DOWNLOAD RECORDS FOR PARENT OBJECT INSTANCES
    FOR EACH RELATEDDOCKINGOBJECT LOOP
50 RUN PARAMETERSQL TO GET NEWPRIMARYID OF RELATEDDOCKINGOBJECTS
    FOR EACH NEWPRIMARYID RETRIEVED LOOP
        -- CHECK IF THE INSTANCE OF THE OBJECT IS VISIBLE TO THE LAPTOP
        USER
        CHECKOBJECTVISIBILITY (LAPTOPNODEID, OBJECTNAME, PRIMARYROWID)
55 IF VISIBLE THEN
        DOWNLOADOBJECTINSTANCE (LAPTOPNODEID,
                                RELATEDDOCKINGOBJECT,
                                NEWPRIMARYROWID);
        END IF;
60 END LOOP;
    END LOOP;

    RETURN TRUE;
65 }

```

```
-- REMOVE AN OBJECT INSTANCE TO A LAPTOP
-- THIS FUNCTION ALSO REMOVES ALL RELATED DOCKING OBJECT INSTANCES.
BOOL REMOVEOBJECTINSTANCE (LAPTOPNODEID, OBJECTNAME, PRIMARYROWID)
5 {
  -- CHECK IF THE OBJECT INSTANCE IS ALREADY DOWNLOADED TO THE LAPTOP
  FIND THE OBJECT INSTANCE IN THE S_DOBJ_INST TABLE
  IF DOES NOT EXIST ON LAPTOP THEN
    RETURN TRUE;
10  END IF;

  -- DELETE THE OBJECT INSTANCE FROM S_DOBJ_INST TABLE

  -- WRITE REMOVE OBJECT RECORDS TO THE TXN LOG
  FOR EACH MEMBER TABLE OF THE DOCKING OBJECT LOOP
15  GENERATE SQL SELECT STATEMENT TO GET RECORDS TO DELETE
    WRITE EACH RETRIEVED RECORD TO THE USER TXN LOG FILE
  END LOOP;

  -- REMOVE FOR PARENT OBJECT INSTANCES
20  FOR EACH RELATEDDOCKINGOBJECT LOOP
    RUN PARAMETERSQL TO GET NEWPRIMARYID OF RELATEDDOCKINGOBJECTS
    FOR EACH NEWPRIMARYID RETRIEVED LOOP
      -- CHECK IF THE INSTANCE OF THE OBJECT IS VISIBLE TO THE LAPTOP
25  USER CHECKOBJECTVISIBILITY (LAPTOPNODEID, OBJECTNAME, PRIMARYROWID)
        IF NOT VISIBLE THEN
          REMOVEOBJECTINSTANCE (LAPTOPNODEID,
                                RELATEDDOCKINGOBJECT, NEWPRIMARYROWID);
30        END IF;
      END LOOP;
    END LOOP;

  RETURN TRUE;
35 }
```

WE CLAIM:

1. A METHOD OF ESTABLISHING AND MAINTAINING A SECURE TCP/IP SESSION
BETWEEN A SERVER (301) HAVING A DATABASE AND A CLIENT (311),
5 COMPRISING:
 - (A) SENDING A HELLO MESSAGE FROM THE CLIENT (311) TO THE SERVER
(301) ESTABLISHING CLIENT AUTHORIZATION;
 - (B) SAID SERVER (301) SENDING A SERVER SESSION ID AND A SESSION
10 PRIVATE KEY TO THE CLIENT (311).
2. THE METHOD OF CLAIM 1 WHEREIN THE CLIENT'S (311) HELLO MESSAGE
CONTAINS THE CLIENT'S (311) PUBLIC KEY.
3. THE METHOD OF CLAIM 1 COMPRISING ENCRYPTING THE SERVER (301)
15 SESSION ID USING THE CLIENT'S (311) PUBLIC KEY.
4. THE METHOD OF CLAIM 1 COMPRISING SENDING ALL MESSAGES BETWEEN THE
CLIENT (311) AND THE SERVER (301) WITH THE SESSION ID INCLUDED
THEREIN.
- 20 5. THE METHOD OF CLAIM WHEREIN THE SESSION ID IS A UNIQUE, NON-
REPEATING SESSION ID.
6. THE METHOD OF CLAIM 1 WHEREIN EACH MESSAGE IN THE SESSION HAS AN
25 IDENTIFIER.
7. THE METHOD OF CLAIM 6 WHEREIN THE IDENTIFIER IS MONTONICALLY
INCREASING.
- 30 8. THE METHOD OF CLAIM 1 COMPRISING STORING THE CLIENT (311) PRIVATE
KEY AT THE CLIENT (311),
9. THE METHOD OF CLAIM 1 WHEREIN INDIVIDUAL MESSAGES COMPRISE SQL
COMMANDS.
- 35 10. AN ARTICLE OF MANUFACTURE COMPRISING:

A COMPUTER USABLE MEDIUM HAVING COMPUTER READABLE PROGRAM
CODE MEANS EMBODIED THEREIN FOR ESTABLISHING AND MAINTAINING
40 A SECURE TCP/IP SESSION, THE COMPUTER READABLE PROGRAM MEANS
IN SAID ARTICLE OF MANUFACTURE COMPRISING:
 - (A) COMPUTER READABLE PROGRAM CODE MEANS FOR CAUSING A COMPUTER
TO EFFECT SENDING A HELLO MESSAGE FROM THE CLIENT (311) TO
THE SERVER (301) ESTABLISHING CLIENT AUTHORIZATION; AND
 - 45 (B) COMPUTER READABLE PROGRAM CODE MEANS FOR CAUSING THE

COMPUTER TO EFFECT SAID SERVER (301) SENDING A SERVER SESSION ID AND A SESSION PRIVATE KEY TO THE CLIENT (311).

- 5 11. THE ARTICLE OF MANUFACTURE OF CLAIM 10 COMPRISING COMPUTER READABLE PROGRAM CODE MEANS FOR CAUSING THE COMPUTER TO INCORPORATE THE CLIENT'S (311) PUBLIC KEY IN THE CLIENT'S (311) HELLO MESSAGE.
- 10 12. THE ARTICLE OF MANUFACTURE CLAIM 10 COMPRISING COMPUTER READABLE PROGRAM CODE MEANS FOR ENCRYPTING THE SERVER (301) SESSION ID USING THE CLIENT'S (311) PUBLIC KEY.
- 15 13. THE ARTICLE OF MANUFACTURE OF CLAIM 10 COMPRISING COMPUTER READABLE PROGRAM CODE FOR SENDING ALL MESSAGES BETWEEN THE CLIENT (311) AND THE SERVER (301) WITH THE SESSION ID INCLUDED THEREIN.
- 20 14. THE ARTICLE OF MANUFACTURE OF CLAIM 10 COMPRISING COMPUTER READABLE PROGRAM CODE FOR GENERATING IS A UNIQUE, NON-REPEATING SESSION ID.
- 25 15. THE ARTICLE OF MANUFACTURE OF CLAIM 10 COMPRISING COMPUTER READABLE PROGRAM CODE FOR GENERATING AN IDENTIFIER FOR EACH MESSAGE IN THE SESSION.
- 30 17. THE ARTICLE OF MANUFACTURE COMPRISING COMPUTER READABLE PROGRAM CODE FOR STORING THE CLIENT (311) PRIVATE KEY AT THE CLIENT (311),
- 35 18. THE ARTICLE OF MANUFACTURE OF CLAIM 10 COMPRISING COMPUTER READABLE PROGRAM CODE FOR GENERATING SQL COMMANDS AS INDIVIDUAL MESSAGES IN THE SESSION.
- 40 19. A PROGRAM STORAGE DEVICE READABLE BY A MACHINE, TANGIBLY EMBODYING A PROGRAM OF INSTRUCTIONS EXECUTABLE BY A MACHINE TO PERFORM METHOD STEPS FOR ESTABLISHING AND MAINTAINING A SECURE TCP/IP SESSION BETWEEN A SERVER (301) HAVING A DATABASE AND A CLIENT (311), SAID METHOD STEPS COMPRISING:
- 45 (A) SENDING A HELLO MESSAGE FROM THE CLIENT (311) TO THE SERVER (301) ESTABLISHING CLIENT AUTHORIZATION;
- (B) SAID SERVER (301) SENDING A SERVER SESSION ID AND A SESSION

PRIVATE KEY TO THE CLIENT (311).

20. THE PROGRAM STORAGE DEVICE OF CLAIM 19 WHEREIN SAID METHOD
STEPS FURTHER COMPRISE THE CLIENT'S (311) HELLO MESSAGE
5 CONTAINS THE CLIENT'S (311) PUBLIC KEY.
21. THE PROGRAM STORAGE DEVICE OF CLAIM 19 WHEREIN SAID METHOD
STEPS FURTHER COMPRISE ENCRYPTING THE SERVER (301) SESSION
ID USING THE CLIENT'S (311) PUBLIC KEY.
10
22. THE PROGRAM STORAGE DEVICE OF CLAIM 19 WHEREIN SAID METHOD
STEPS FURTHER COMPRISE SENDING ALL MESSAGES BETWEEN THE
CLIENT (311) AND THE SERVER (301) WITH THE SESSION ID
INCLUDED THEREIN.
15
23. THE PROGRAM STORAGE DEVICE OF CLAIM 22 WHEREIN SAID METHOD
STEPS FURTHER COMPRISE GENERATING A UNIQUE, NON-REPEATING
SESSION ID.
- 20 24. THE PROGRAM STORAGE DEVICE OF CLAIM 19 WHEREIN SAID METHOD
STEPS FURTHER COMPRISE ASSIGNING EACH MESSAGE IN THE SESSION
AN IDENTIFIER.
- 25 25. THE PROGRAM STORAGE DEVICE OF CLAIM 24 WHEREIN SAID METHOD
STEPS FURTHER COMPRISE ASSIGNING EACH MESSAGE IN THE SESSION
A MONTONICALLY INCREASING IDENTIFIER.
26. THE PROGRAM STORAGE DEVICE OF CLAIM 19 WHEREIN SAID METHOD
STEPS FURTHER COMPRISE STORING THE CLIENT (311) PRIVATE KEY
AT THE CLIENT (311),
30
27. THE PROGRAM STORAGE DEVICE OF CLAIM 19 WHEREIN SAID METHOD
STEPS FURTHER COMPRISE INCORPORATING SQL COMMANDS IN
INDIVIDUAL MESSAGES.

1 / 9


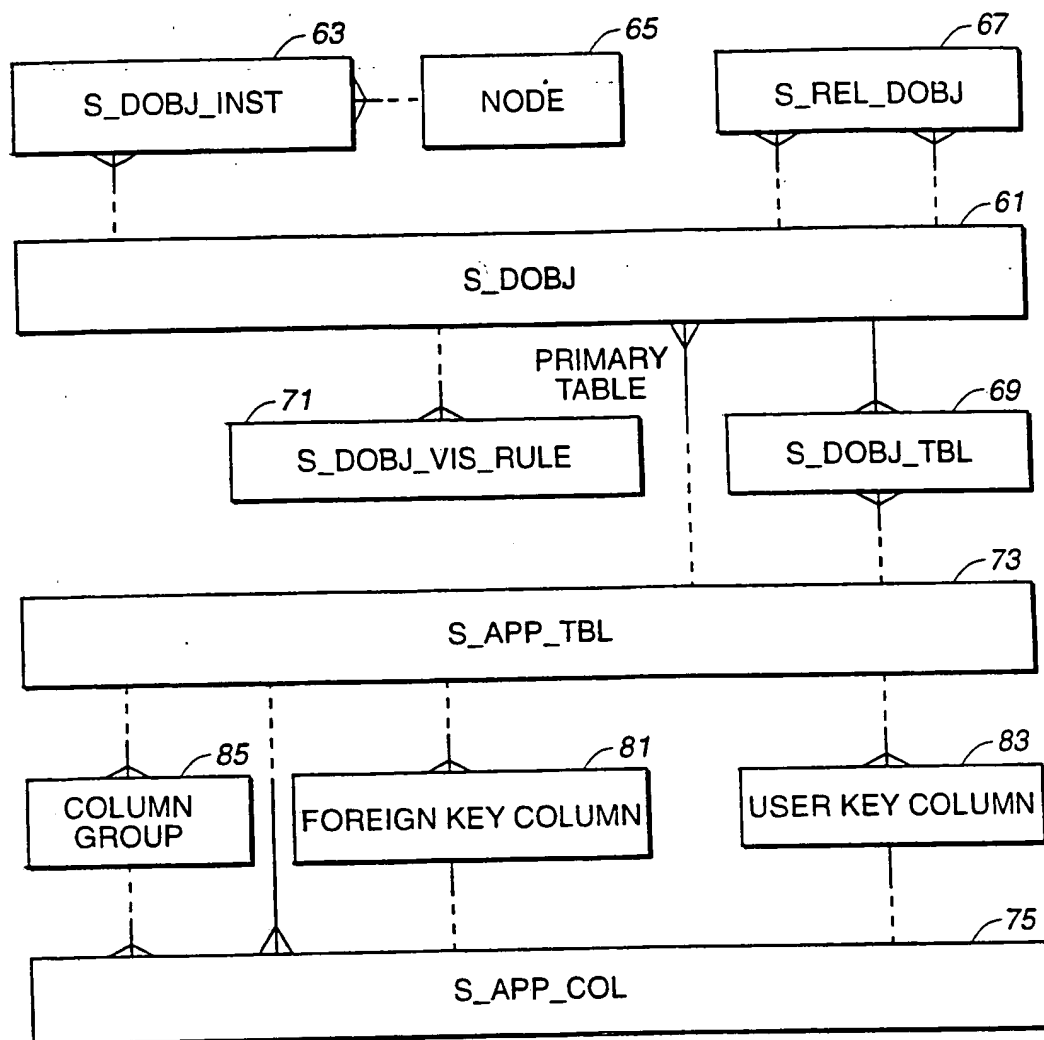



FIG. 1


2 / 9

**FIG. 2**

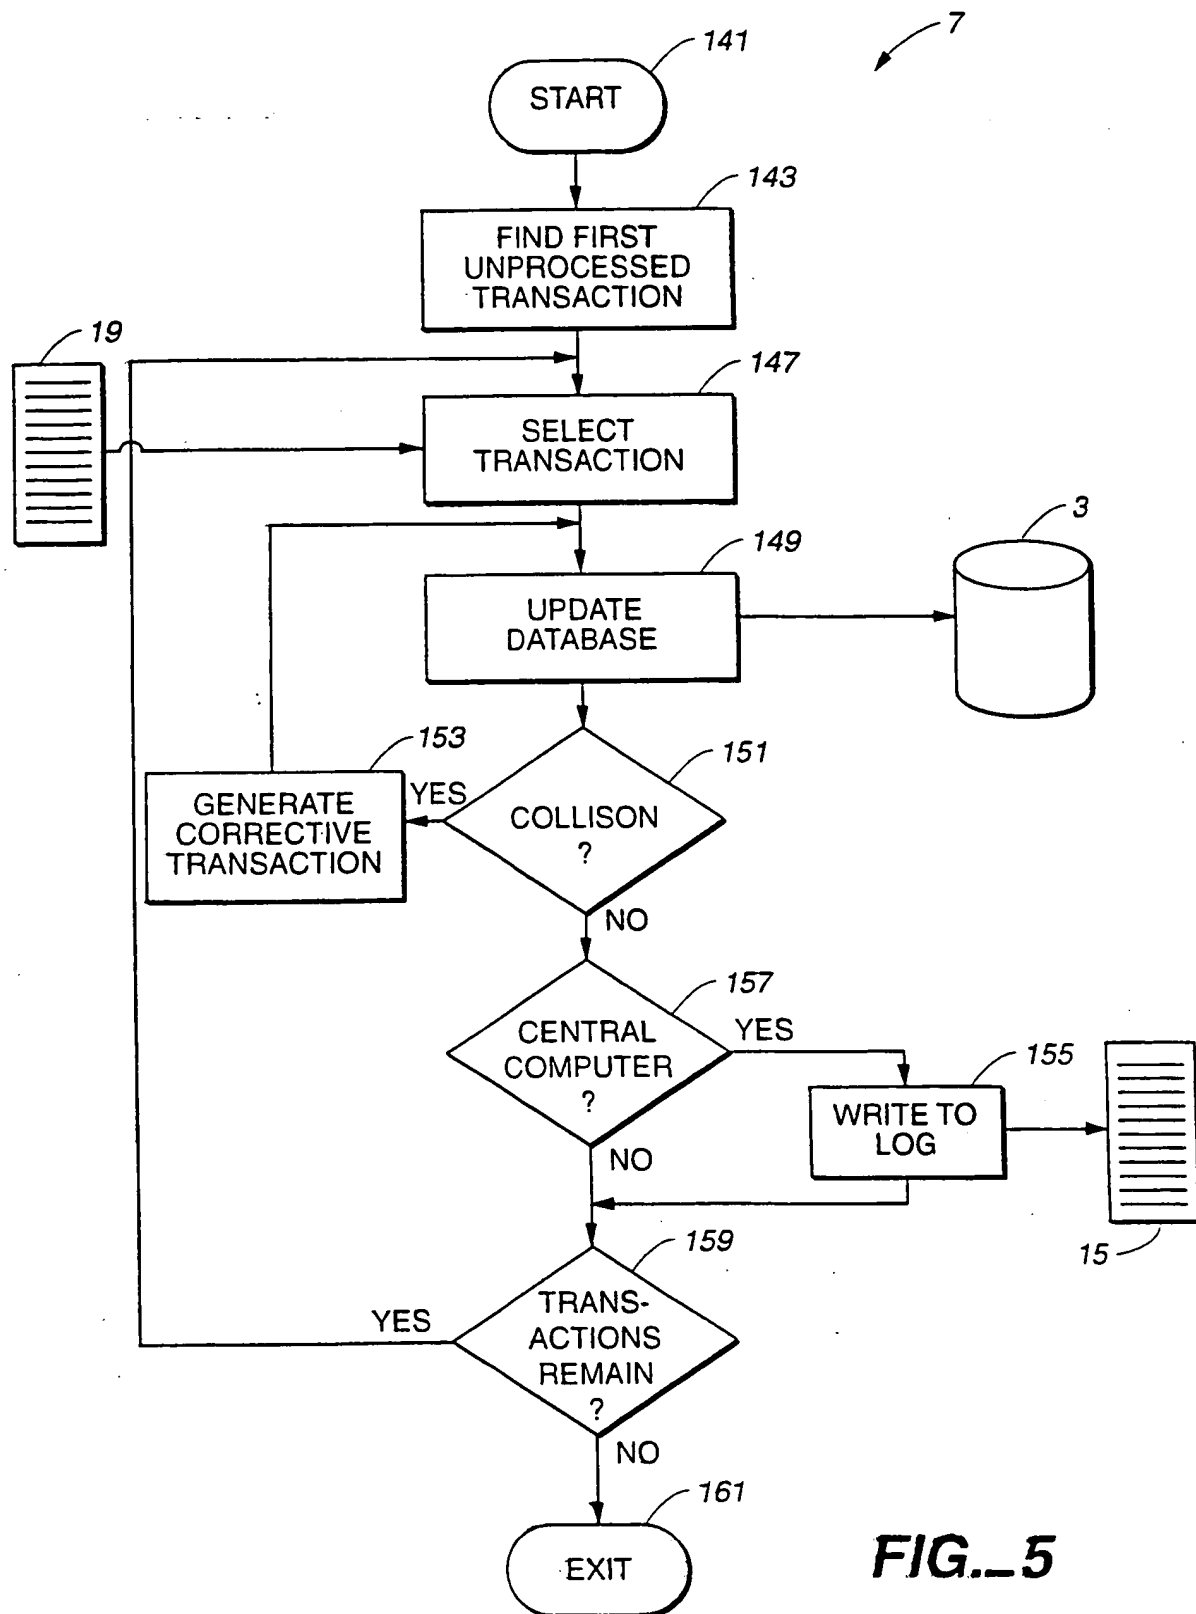
3 / 9

**FIG. 3**

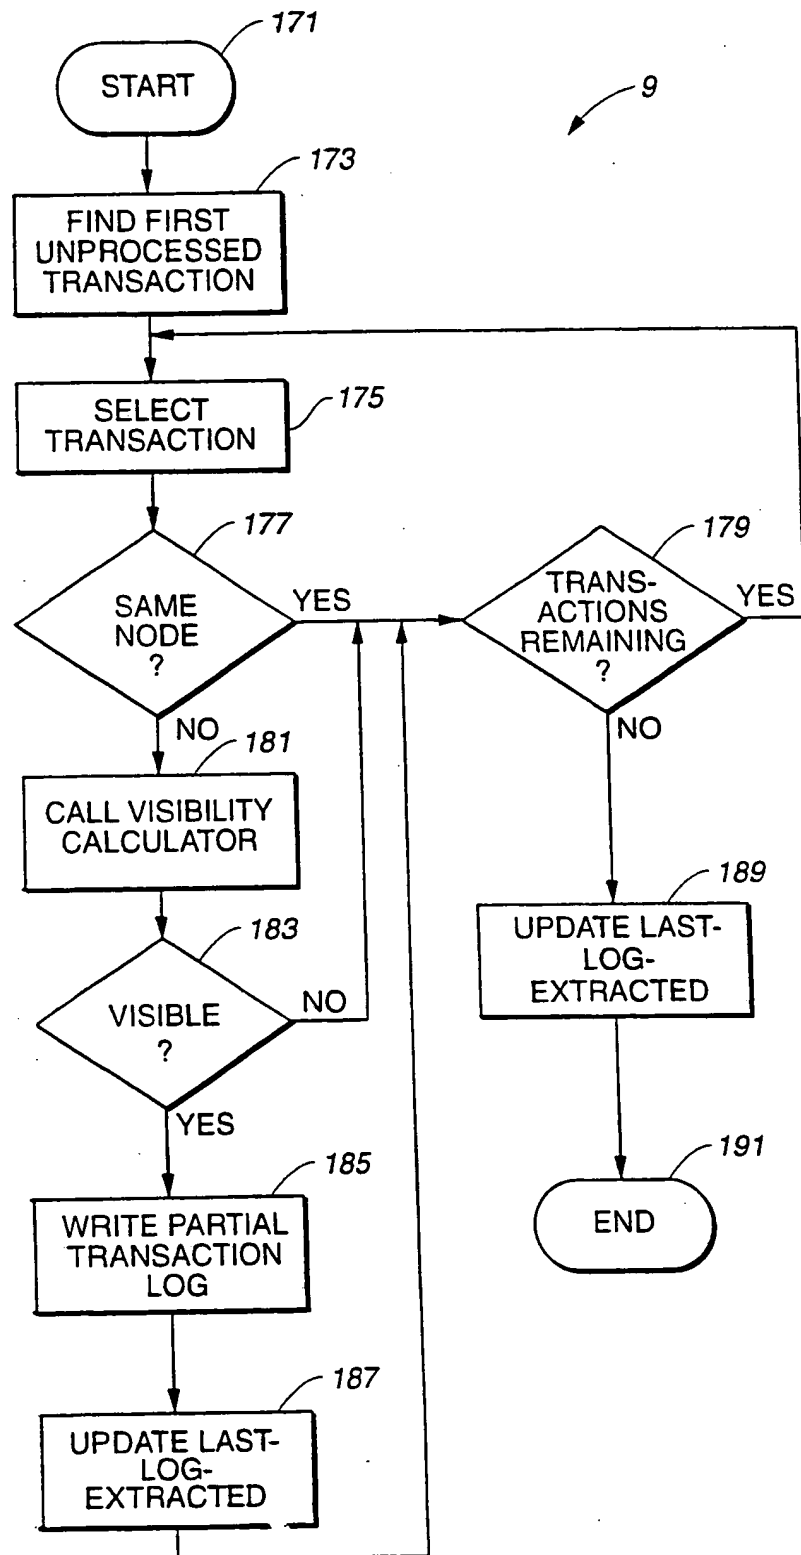
4 / 9

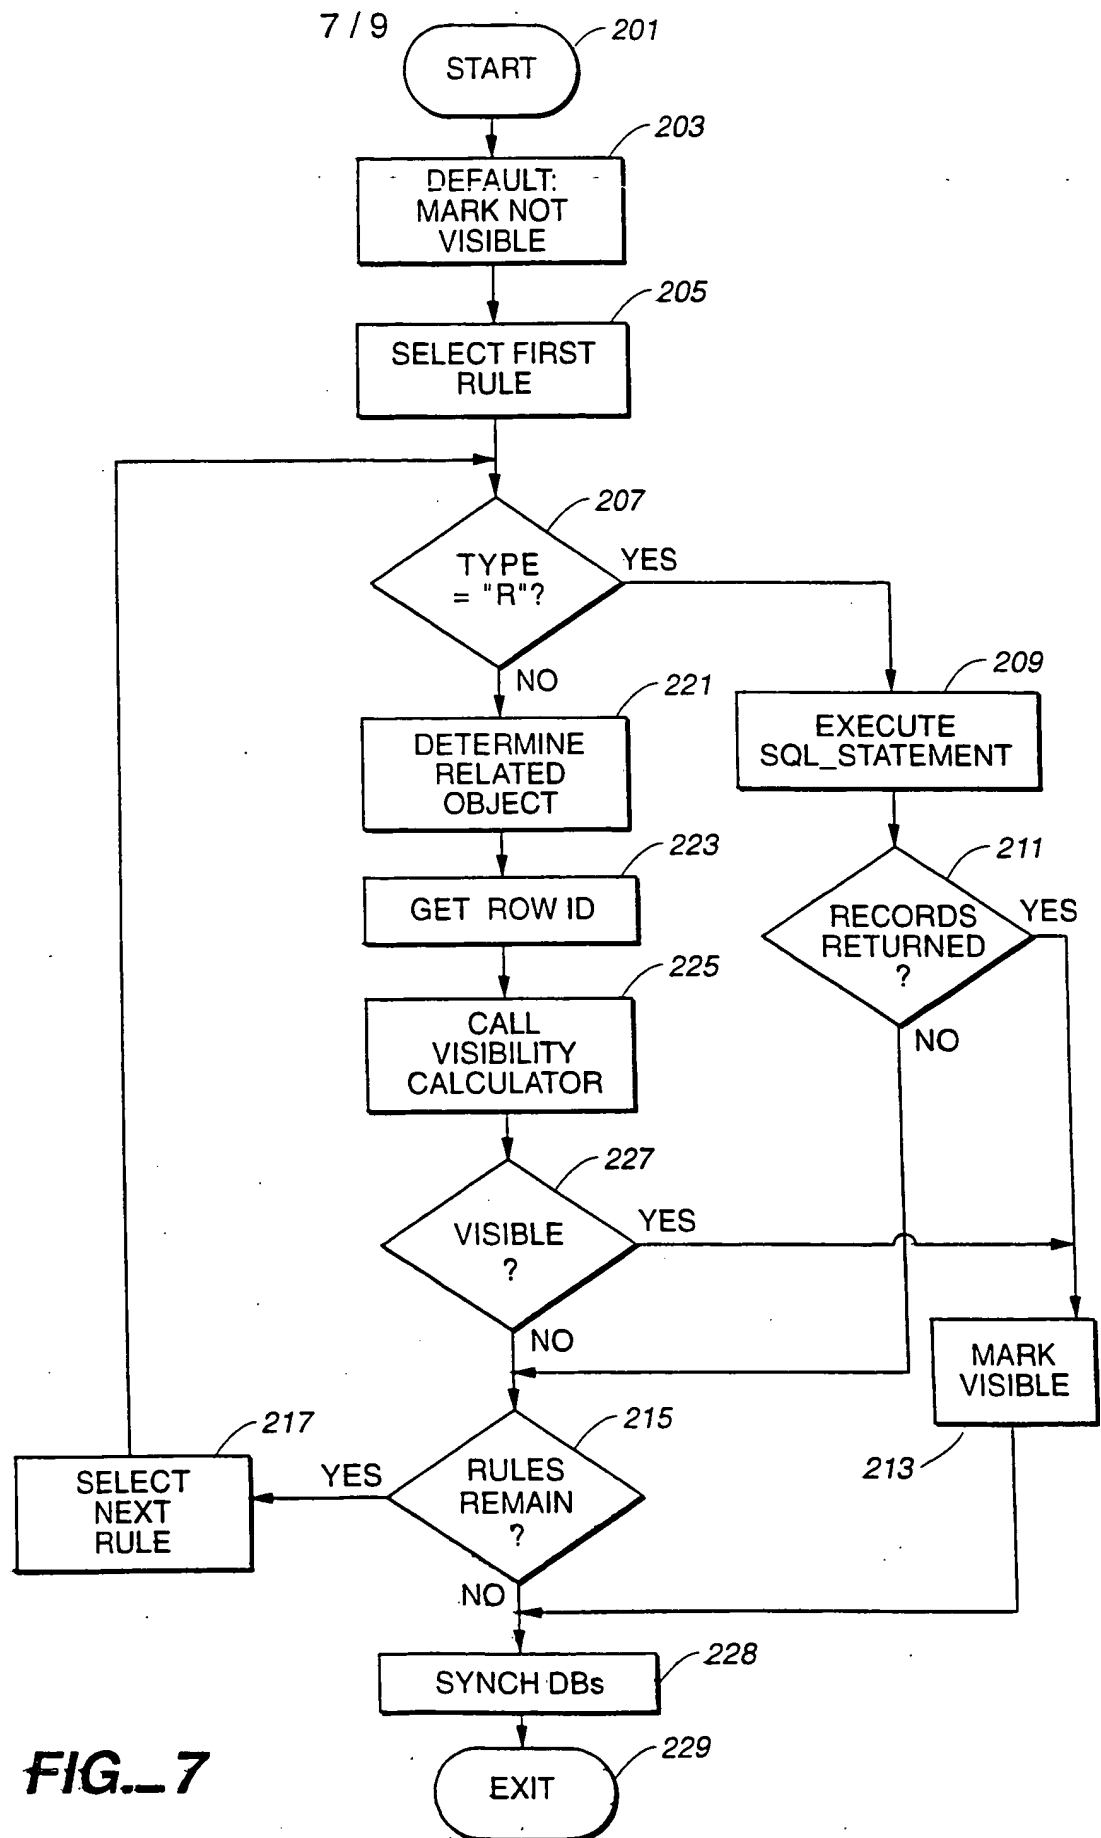
**FIG. 4**

5 / 9

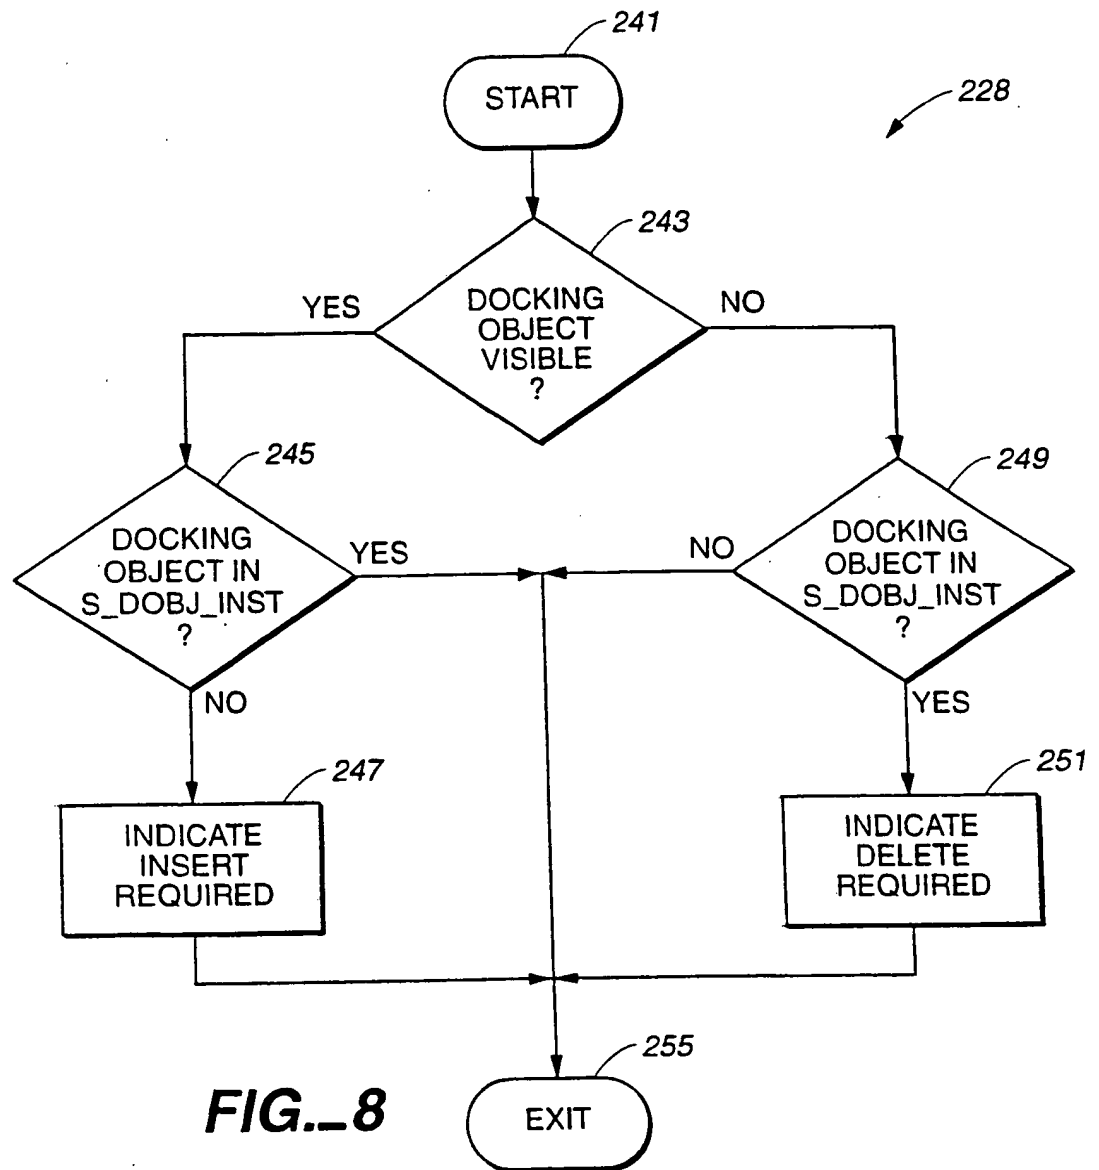
**FIG. 5**

6 / 9

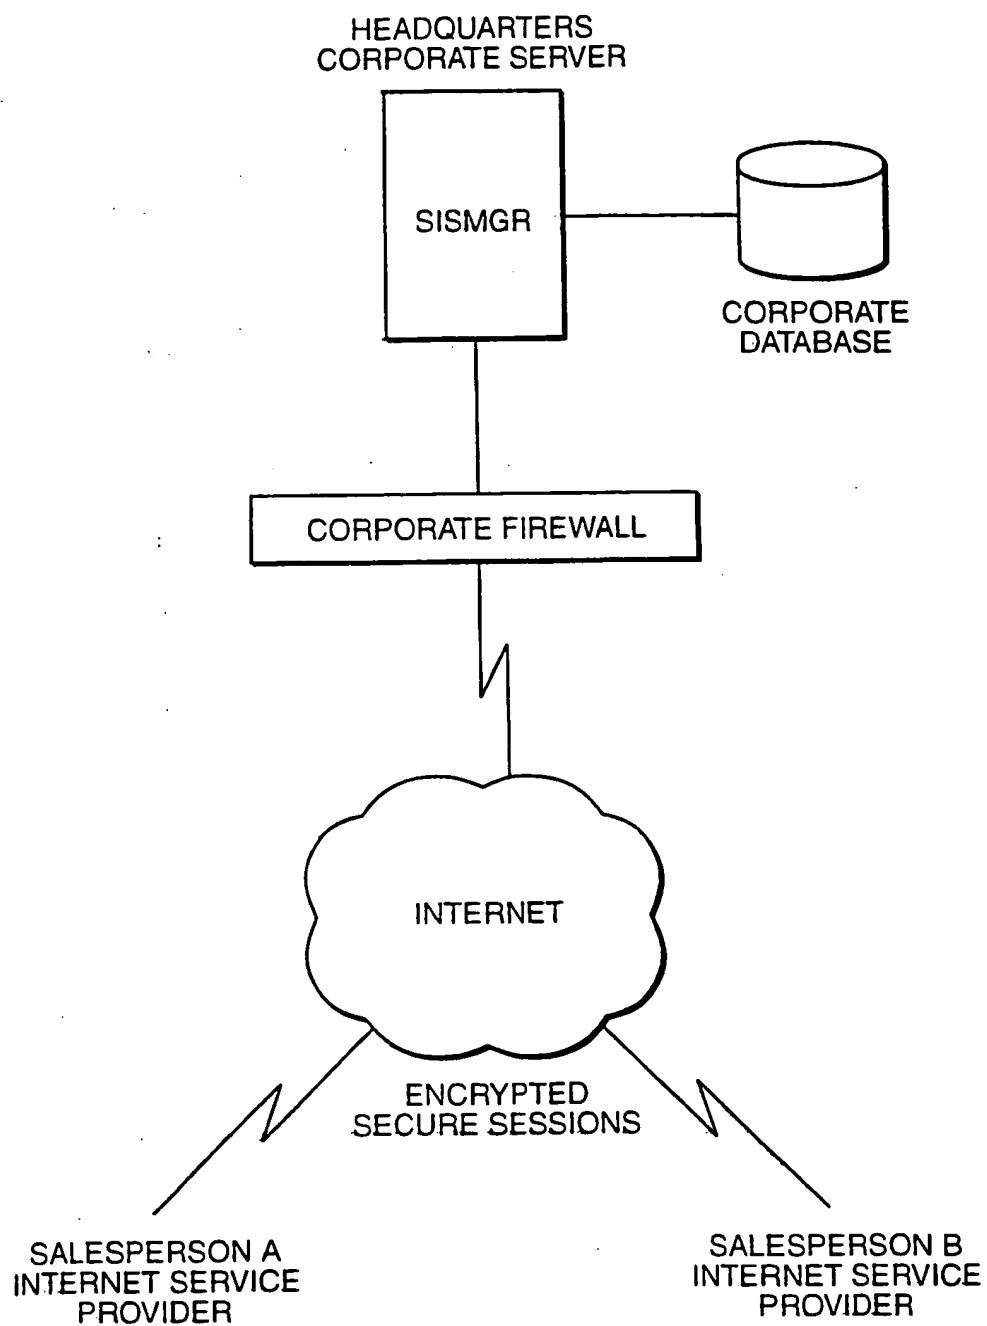
**FIG. 6**

**FIG. 7**

8 / 9



9 / 9

**FIG. 9**

THIS PAGE BLANK (USPTO)